Institut d'études politiques de Paris
ÉCOLE DOCTORALE DE SCIENCES PO
Programme doctoral en économie
Département d'Économie
Doctorat en sciences économiques

Returns to Education and Experience on the Labor Market

A Matching Perspective

Pauline Corblet

Thesis supervised by
Alfred Galichon, Professor of Economics, New York University

To be defended on May 18, 2022

Jury:
Ms. Zsofia Bárány, Associate Professor of Economics, Central European University
Mr. Moshe Buchinsky, Professeur Titulaire, Sciences Po Paris
Mr. Arnaud Dupuy, Full Professor of Economics, University of Luxembourg
Mr. Jeremy Fox, Professor of Economics, Rice University
Mr. Alfred Galichon, Professor of Economics, New York University
Mr. Jean-Marc Robin, Professeur des Universités, Sciences Po Paris

Abstract

This thesis comprises three self-contained chapters. Its main object of interest is the wage returns to education and experience on labor markets, and the earning inequalities these returns generate. To understand the origin of returns to education and experience, and why they vary across time and space, this thesis adopts a matching perspective: it investigates the determinants of relationship formation between workers/employees and firms/employers. Employer-employee relations rest on two elements: first, the profitability of the relation, i.e. how much a worker and firm can produce together, and second, relative supply and demand: how many workers are available to work in firms and vice-versa. Because they determine which relationships are formed, these two elements eventually impact wages paid by firms to workers. The first chapter documents flattening wage returns to experience between higher education graduates entering the French labor market in 1998 and 2010. I decompose differences in average wage growth by occupation into an extensive (a composition effect across occupations) and intensive margin (a variation in wage growth within occupations). I then study two mechanisms behind the wage growth slow down: access to managerial positions and impact of initial match quality. I find access to managerial positions is more infrequent for recent cohorts. I also find that initial match quality has not worsened between the 1998 and 2010 cohorts, but its impact on future wages has become more enduring. The second chapter studies the interplay between worker supply and firm demand, and their effect on sorting and wages in the labor market. Specifically, I investigate a decrease in the education wage premium on the Portuguese labor market between 1987 and 2017. I build a model of one-to-many matching with multidimensional types in which several workers are employed by a single firm. I structurally estimate the model on matched employer-employee data. Counterfactual exercises suggest that both changes in worker preferences and the increasing relative productivity of high school graduates over non-graduates act as a mitigating force on the decreasing high school wage premium, but do not fully compensate for high school graduates' rise in relative supply. In the third chapter, co-authored with Jeremy Fox and Alfred Galichon, we explore how expectations on future returns influence matching decisions. We introduce a model of dynamic matching with transferable utility. We explore aggregate dynamics and show that a stationary equilibrium exists. We propose two algorithms to compute a stationary equilibrium, and adapt both for estimation. We then use the methods developed to estimate a model of geographic mobility costs for Swedish engineers. We find that mobility costs impose a sizeable penalty in match production, and evolve non-linearly by age.

Acknowledgements

My most grateful thanks go to my advisor Alfred Galichon. I first met Alfred when I became his research assistant during my masters in Sciences Po, and was immediately struck by his intelligence and enthusiasm. Alfred showed me how powerful a tool convex optimization could be in economics, and has inspired me to pursue a PhD under his supervision. Throughout my PhD, he has been very generous with his time and has provided me with invaluable guidance and support, both on general and technical subjects. I am very much indebted to him.

I am also very grateful to my advisors Jean-Marc Robin and Moshe Buchinsky. They have both taken the time to follow my work and give me valued feedback and advice throughout my PhD, from which my research has substantially benefitted. Both their support has been especially precious during my job market.

I would like to thank Zsofia Bárány and Kerstin Holzheu for their constant encouragements and feedbacks. This dissertation owes a lot to their expertise in Labor economics and Macroeconomics. I also want to thank Pierre Cahuc for taking the time to read my job market paper, which has been tremendously helpful, and Florian Oswald for providing me with insightful feedback and encouraging me to hone and improve my computational skills.

I would also like to take this opportunity to thank Jeremy Fox and Arnaud Dupuy who, together with Jean-Marc, Moshe and Zsofia, have agreed to be part of my jury. I am very happy to have worked with Jeremy as a co-author: I have learned so much at his contact. I very much look forward to working with Arnaud, as I start a postdoc at the University of Luxembourg next year.

My work has also benefitted from the many perceptive comments of the faculty members at the Department. In particular, I would like to thank Ghazala Azmat, Nicolas Coeurdacier, Jeanne Commault, Michele Fioretti, Guy Laroque, Marleen Mara, Isabelle Méjean, Stefan Pollinger, Thierry Mayer and Xavier Ragot.

I would also like to thank the administrative staff at Sciences Po: Florence Avisse, Cathy Benard, Pilar Calvo, Séverine Germain, Claudine Lamaze, Melissa Mundell, and Hadjila Nezlioui-Serraz, thank you for you availability and efficiency in dealing administrative issues. Special thanks go to Sandrine Le Goff for overseeing traveling and accommodation, always finding an available room when we need it, and organizing the Lunch Seminar! Thanks also go to Guillaume Sarrat de Tramezaigues, for making the practical side of this PhD a smooth experience.

I have had a wonderful time at Sciences Po throughout my PhD thanks to its many smart and friendly PhD students: my own PhD cohort, Edgard, Daniel, Jérôme, Pierre, Stefan Victor, and Zydney; my job market partners: Alaïs, Ariane, Clément, Julia, Oliver

and Vladimir; and everyone who made our open space *the combles* a great place to work: Alexis, Clémence, Dàniel, Gustave, Juan, Leonard, Marco, Moritz, Mylène, Naomi, Nourhan, Riddhi, Samuel, Ségal, Valentin and Victor. There was always someone in this joyous crowd ready for a brainstorming or a drink, and it has greatly contributed to the completion of this thesis. I wish everyone all the best. I would also like to thank the now-doctors who defended their thesis during my time here, their advice early on in my PhD and their career success have been most inspirational: Aseem, Camille, Charles, Elisa, Edoardo, Julien, Ludovic, Nicolò, Sophie and Tyler.

Writing a doctoral thesis can be a nerve-racking experience, and I am grateful for all my brilliant friends who reminded me daily that life is not only about work: Howon, Grégory, Thomas, Christina, Henrik, Anita, Léa and all the 'Scubes' with whom I started my education at Sciences Po more than ten years ago; Mélina, Patrick, Cécile, Rancy, Léa, Max and Guillermo who can always be counted on to vent about their own PhDs around a good dinner; Laurène, Robin, Kevin and John, who do not despair of improving my theatrical and cinematographic culture; Alice, Florestan, Clotilde, Sophie, Laurence, Manon and Manon who are always ready for a good laugh when I visit Belgium; and the friends that live further away but are not less dear: Flora, Laura, Pierre, François, Arnaud and Natalie.

I would also like to thank my parents, Sandrine and Charles, and my brother, Alexandre, for their constant love, support and interest in my work. My mother predicted I would be a researcher when I was five, and as always she was right. I am also indebted to my grand-parents, Patrice and Gérard, my aunt and uncle, Ariane and Antoine, and my cousins Danaé, Olympe and Hector, for welcoming me with open arms when I moved to Paris eleven years ago, and caring for me since. I also thank my cousins in La Réunion for continuously inviting us for much needed vacations in the sun, and my family-in-law for introducing me to the joys of cross-country skiing. I am also grateful for the encouragements and love of my grand-mother Monique and grand-father Jean-Claude, who is dearly missed.

Finally, I would never have completed this thesis without the unbending support and faith of my partner Thibaut. He has lead the way into the PhD venture, and has helped me progress in so many ways since: by tipping me on algorithms, constantly refactoring my code, and proofreading this thesis, to only cite a few. I am extremely lucky to have him in my life, and I cannot wait to start the new phase of our life together.

Declaration

My work has received the support of the European Research Council (ERC Consolidator Grant no. 866274).

Contents

List of Figures

List of Tables

Introduction

Labor markets are a pervasive aspect of 21^{st} century Western societies. Every single individual in Europe and the United States has at some point in their life an experience on labor markets, either by looking for employment or being employed, and scarcely a day goes by without media and governments scrutinizing and dissecting labor trends. Of particular interest are the unemployment rate and wage levels, which are both viewed as informative health measures of the economy. In the academic field, a large branch of economic research is also dedicated to labor markets, and questions pertaining to the impact of trade and globalization, the workforce's educational composition, the setting of wages, and the role of labor market institutions (among others) provide lively research areas.

This dissertation aims at understanding labor markets from a microeconomics point of view: it models both workers' and firms' individual choices on the labor market as a utility (for workers) or profits (for firms) maximizing, which at the aggregate level translates into workers' supply of labor and firms' demand for labor. It then uses the supply and demand framework to explain the drivers of wage setting. As such, it falls within the established tradition of neoclassical labor economics, in which wages are equilibrium quantities determined by supply and demand. However, it departs from the baseline theory as presented by Hicks (1932): by accounting for individual idiosyncrasies, introducing heterogeneous workers and firms, and recognizing imperfect competition on labor markets, this dissertation reckons with a richer wage-setting framework than the standard view that wages are simply equal to the marginal product of labor.

In this dissertation, I choose to focus on an important aspect of labor markets, namely wage returns to education and experience. It is generally observed that workers who have attained higher education levels, such as high school or university graduates, are paid higher wages than workers who have not. Workers who have more years of experience in the labor market are typically also paid more than their less-experienced peers. Returns to education and experience often combine, in favor of educated and experienced workers. However, the

wage gaps between educated and uneducated (or less educated), and experienced and inexperienced workers vary across time and space. Understanding why these gaps exist in the first place, and why they differ by decade and region is essential for two reasons. First, because they are important drivers of inequality: in 2016 wages accounted for 55% of total household income in France, and 72% in the United States (Rani and Furrer (2016)). Because wages are the main source of total household income in most countries, wage inequality across households translates into income inequality (Autor (2014)). This is especially true among the first 90$^{\text{th}}$ percentile of the income distribution: capital income is concentrated at the top of the income distribution and is the main driver of inequality between the poorest 90% and the richest 10% (Krueger et al. (2010)). The second reason why understanding wage gaps is necessary is that they inform us on how firms function. Indeed, how much firms agree to pay different types of workers depends on their needs: all other things equal, a firm is ready to pay much more a worker that possesses a set of skills crucial to its functioning than a worker who does not. The increase in the college wage premium (i.e. the average wage gap between workers who graduated from college and those who did not) between the late 1970's and early 2000's in the United States, United Kingdom, and Canada as evidenced in Krueger et al. (2010), is convincingly interpreted by a large literature as the consequence of *skill-biased technological change*: a shift in firms' production structure that made educated workers more productive relative to uneducated workers (Katz and Murphy (1992), Acemoglu (1998), Goldin and Katz (2008), Autor et al. (2020)). More recently, the college wage premium trend has flattened, and it has even decreased in some European countries (for instance in Germany, Italy, Spain). This shift in trend is indicative of broad changes in the labor markets where it has occurred and could be either due to a change in demand that would go against skill-biased technological change or to a change in supply: both Europe and the United States have experienced large education expansions over the past 50 years. If supply changes are at play, even if firms' demand for educated workers has persisted, it could be that their larger number is causing the decrease in relative wage. Understanding how supply and demand play out in determining wage gaps is central to economic policy-making and should shape education and production policies [1].

Being able to tell apart demand and supply effects' on wage levels is at the heart of this dissertation. To do so, it develops a set of modeling and data tools that borrow from various branches of the economics literature. First, I use ordinary least square regressions in the

[1] For instance, the French high school curriculum was reformed in-depth in 2019, leading to a strengthened specialization of students. Whether this will prove a strength for new graduates in the labor market remains to be assessed.

first chapter of this dissertation to assess the average impact of one or several explanatory variables on an outcome variable. Second, I rely on the structural econometric literature and the matching literature to build a novel model of matching between workers and firms in the labor market. The model can be either static, as in the second chapter, or dynamic with forward-looking agents, as in the third one. In both cases, the model incorporates workers who maximize their utilities, and firms who maximize their profits, with wages acting as equilibrium transfers. Importantly, both workers and firms are heterogeneous in multiple dimensions: education, age or occupation for workers, sectors or location for firms. Because workers perceive different amenities in different firms, and not all workers are equally productive in firms, the model is able to generate rich matching and wage distributions. Finally, I exploit high-quality matching data with rich information on which workers are employed at which firm and for what wage. Using this data, I am able to structurally estimate the aforementioned models using observed matching and wage distributions, which lets me back out the parameters of worker supply and firm demand.

Related literature. This dissertation belongs to the large field of labor economics, which investigates workers' and firms' outcomes (wage, employment, etc.) on labor markets. It is related to several strands of the literature: first to the broad body of research on the mechanisms of wage setting, second to the theoretical matching literature, and third to the literature studying returns to education, and how they have changed over time and across economies.

Investigating determinants of wage is central to the field of labor economics, and I give here only a partial view of the extensive corpus of research on the topic. The neoclassical view, as described by Hicks (1932), is that in the simple setup where a representative firm seeks to maximize profits by hiring its workforce, workers' wages are equal to their marginal product in the firm. This simple model of wage setting can be augmented by accounting for workers' diverse education or experience levels, as in Katz and Murphy (1992) or Card and Lemieux (2001), or by assuming inelastic labor supply, as Manning (2021) describes. An important alteration of this view, which took place in the 1970s and started the search literature (McCall (1970), Mortensen (1970)), is to acknowledge the existence of frictions in workers' search for a job and firms' search for an employee: a worker does not immediately find a job the moment she enters the labor market, and a firm cannot instantly hire whoever is needed. The search literature provides frictions as a rationale for inelastic labor supply, and they are usually modeled with a measure of randomness in the meeting process of workers and firms. Because agents are forward-looking, these frictions impact the setting of wage: when a worker and a firm meet they account for future search costs should they

refuse the present match. The first model to fully endogenize wage distribution in the search literature is the Mortensen-Diamond-Pissarides model (Mortensen (1982), Diamond (1982), Pissarides (1985)). It does so by decomposing the search into two steps: first workers and firms search for a good match, instead of simply having workers searching for a good wage. Wage is then bargained as a split of wage surplus. Another approach is wage posting, initially developed by Burdett and Mortensen (1998) is on-the-job search: even when they already have a job, workers keep looking for a potentially better one: this creates a job ladder with a full wage distribution. Both the wage bargaining and the wage posting views have since been extended and enriched since: Postel–Vinay and Robin (2002) augment wage posting with price discrimination from firms based on workers' characteristics, Mortensen (2003) examines various wage-setting mechanisms in a model of on-the-job search, Cahuc et al. (2006) augment the wage bargaining model with on-the-job search, accounting for firms' offers and counter-offers. Another important contribution by Shimer and Smith (2000) introduces search frictions in the matching model of Becker (1973) and characterizes the assortativeness pf equilibrium. Shimer and Smith (2000)'s result has spawned a rich sub-literature in search-matching models: Lise et al. (2016) extend it to one-the-job search, Lise and Robin (2017) use the framework to understand how a non-)stationary environment impacts wage setting, Moscarini and Postel-Vinay (2018) emphasize the role of the job ladder created by frictions in amplifying aggregate shocks, Lise and Postel-Vinay (2020) considers multidimensional worker types and firms requirements. Another critical extension to the search model is directed search (Julien et al. (2000), Burdett et al. (2001), Menzio and Shi (2011)) that postulates that workers and firms are able to direct their search based on counter-parties characteristics, or in the case of wage posting models, workers direct their search depending on the wages posted by firms. See Hosios (1990) for an important link between wage bargaining models and efficiency in wage posting models. The directed search set up can be extended to simultaneous job applications, as in Albrecht et al. (2006), Kircher (2009) and Galenianos and Kircher (2009), or to job-to-job transitions (Garibaldi et al. (2016)), idiosyncratic risk (Schaal (2017)) or selective effort (Tsuyuhara (2016)). Directed search model are useful tools to understand wage rigidity (Menzio and Moen (2010)), business cycles (Menzio and Shi (2011)) or sorting (Eeckhout and Kircher (2011)).

Parallel to the search structural literature, another view of wage-setting has developed, that does not focus on search frictions but instead on worker and firm heterogeneity, following the seminal work by Abowd et al. (1999) (hereafter AKM), who empirically estimate the contribution of workers and firms to earnings dispersion. To do so, they use 'mover', workers who switch firms, which allows them to separately identify additive worker and firm fixed effects. They find that firms account for about 20% of the wage dispersion, so that

worker heterogeneity appears to be the main driver of the observed heterogeneity in wage. Following in their steps but with contrasting results, Woodcock (2008), Card et al. (2013) and Song et al. (2019) evidence the role of firm wage premium in the rising wage inequality. Importantly Abowd et al. (1999) rely on an exogenous mobility assumption, that is likely not verified in the data: workers' movement between firms must be uncorrelated with their previous earnings. A large structural literature seeks to abstract from this assumption and also aims at capturing sorting, i.e. the combined effect of worker and firm heterogeneity on wage. It does so mainly by directly modeling sorting, as Hagedorn et al. (2017), Lopes de Melo (2018) and Bagger and Lentz (2019). Bonhomme et al. (2019) provide a comprehensive framework to reconcile the structural and empirical approach, which is used by Lentz et al. (2018) to evidence sorting patterns in Denmark.

This dissertation shares the same goals as the aforementioned literature: to understand how wages are set in the labor market. The tools I used to explore this question are related to but different from the ones just described however: I choose to use matching models to describe labor markets. Instead of focusing on embedded frictions like the search literature does, or unobserved productivity, matching models put the spotlight on supply and demand effect, as well as worker and firm surplus derived from observable characteristics, such as education, experience, occupation or sector.

The matching literature is split into two main strands: matching with transferable utility, for which a comprehensive theory is provided in Shapley and Shubik (1971) and matching with non-transferable utility in the framework of which Gale and Shapley (1962) developed the eponymous algorithm. The latter is well-suited to the analysis of problems such as medical residency Roth (1984) or kidney exchange Ashlagi and Roth (2012). Menzel (2015) shows how to identify and estimate preference parameters on a large matching market with non-transferable utility. The former type of matching, with transferable utility, has first caught the attention of family economist with Becker (1973)'s seminal work on marriage markets. The union of marriage economics and matching theory has spawned a rich literature in family economics: since Choo and Siow (2006) proposed a structural model of matching between marriage partners with transferable utility to estimate partners' utility parameters when marriage patterns are observed, many have followed suit: Dupuy and Galichon (2014) study mutual attraction between personality traits, Chiappori et al. (2017) develop a model with investment on children, Ciscato et al. (2020) compares patterns in homosexual and heterosexual marriages and Chiappori et al. (2020b), Chiappori et al. (2020a) survey the evolution of assortative matching in the US and the UK. Alternative approches to the Choo and Siow (2006)'s view include Chiappori et al. (2012) who simplify the parametrization of surplus to be able to estimate the model through reduced form equations, the approach developed

in Fox (2010b), Fox and Bajari (2013) Fox et al. (2018), that focuses on identifying the distribution of unobserved heterogeneity using several markets, and Sinha (2018), Gualdani and Sinha (2022) who aim their attention at identification of non-parametric models. Recently, the theory of matching with transferable utility has benefited from advances made in optimal transport (Villani (2009), Peyré and Cuturi (2020)). The link between matching and optimal transport is explicited in Galichon (2016) and Galichon (2021), and rests on Gretsky et al. (1992)'s duality proof. It can be summed up as a first welfare theorem: under a large market assumption, the equilibrium resulting from decentralized matching market with idiosyncratic heterogeneity is the same as the equilibrium obtained through solving a social planner problem that maximizes expected surplus with a penalty due to the idiosyncratic heterogeneity. It turns out this social planner problem is nothing different from a regularized optimal transport problem, and can be solved using techniques from this literature (in most cases, the Sinkhorn algorithm). Also see Galichon and Salanié (2021) for a general exposition. On a related topic, the equivalence of matching markets and hedonic markets in the spirit of Ekeland et al. (2004) has been made by Chiappori et al. (2010), and the equivalence between matching and discrete choice à la Berry et al. (1995) is established in Bonnet et al. (2015). Matching with transferable utilities in the spirit of Choo and Siow (2006) is also used to understand labor markets: Dupuy and Galichon (2022) show how using matching and wage distribution separately identifies worker amenities and firm production. Dupuy et al. (2020) introduce a matching market with imperfectly transferable utility to account for taxation, Galichon and Hsieh (2018) develops a matching model with waiting queues, bringing the theory closer to search models, and Dupuy et al. (2021) study the job market of CEOs in Denmark. Another strand of the literature directly combines matching and search to understand sorting (Eeckhout and Kircher (2011)), building on the seminal work of Shimer and Smith (2000). Finally, labor markets are better understood through the prism of one-to-many matching, since the most common form of organization for a firm is to hire several workers. The first step towards one-to-many markets was first completed by Kelso and Crawford (1982) who propose a matching algorithm under the condition that workers are gross substitutes and Hatfield and Milgrom (2005) who develop matching with contracts with substitutable workers. Further development of one-to-many matching models was faced with a sizeable issue: without the restrictive gross-substitute condition, assignment stability could not be guaranteed. The problem has also been explored in auction theory by Bikhchandani and Ostroy (2002) and Vohra (2011), and has been recently solved for matching models through a large market assumption that allows for worker complementarity, see the recent survey by Azevedo and Hatfield (2018), and Che et al. (2019) on non-transferable one-to-many matching with worker complementarity. Recently, there has been a surge in

interest in how matching can explain labor market trends, for instance Lindenlaub (2017) explicit sorting on manual and cognitive skills, Choné and Kramarz (2021) study outsourcing of workers' skills, and Boerma et al. (2021) examine a generalization of assortative matching on one-to-many markets.

Finally, this dissertation is connected to the vast literature on the education wage premium, that analyzes wage gaps between educated and uneducated workers. Many different angles and approaches constitute this literature, and I presently focus on two: the changes in skill/education wage premium, explained through skill-biased technological change, and life-cycle earning profiles, that focus on cumulative returns to education and experiences. The skill-biased technological change literature originally aims at explaining the increase in relative wages of skilled, or college-educated workers compared to their unskilled peers that occurred in the 1970s and 1980s in the United States. Its core idea is that production technology is changing at the time in a way that makes skilled workers increasingly more productive than unskilled workers, which, according to neoclassical wage setting, increases their wage compared to unskilled workers. Skill-biased technological change (SBTC) is first evidenced through estimation of the neoclassical firm production model on wage data by Bound and Johnson (1992), Katz and Murphy (1992), and then grounded into the growing computerization of the US economy (Krueger (1993), Acemoglu (1998), Autor et al. (1998)) and capital-skill complementarity Krusell et al. (2000). Acemoglu (2003), Bloom et al. (2016) also show how trade can induce skill-biased technological change. Writing on the consequences of SBTC Autor and Dorn (2013) show it not only stretches the wage structure between skilled and unskilled but also causes the polarization of the US economy, whereby low and high skilled occupations grow, while middle-skilled occupations decline. Polarization has also been shown to be caused by offshoring (Goos et al. (2014)) and structural change (Bárány and Siegel (2018)). After the wage structure stretch of the 1970s and 1980s, the education wage premium has flattened in the 1990s, which the SBTC hypothesis struggles to explain since computers continued to gain ground in that decade (Card and DiNardo (2002)). In European countries such as France (Verdugo (2014)) starting in the 1970s, the UK (Blundell et al. (2022)) in the 1990s, or Germany (Doepke and Gaetani (2020)) since the 1980s the wage structure has condensed rather than expanded, which is explained through a combination of a supply effect, through education expansions, that overtakes the demand effect from SBTC, and the role of labor market institutions and employment protection, which dampens the impact of SBTC. In the US, the reasons behind the flattening of the college wage premium are still debated, some authors arguing it is due to a decrease in demand for cognitive skills (Beaudry et al. (2015), Valletta (2016)), while others defend on the contrary that demand for cognitive skills has increased in recent years (Blair and Deming (2020))

and that the Great Recession has accelerated SBTC (Hershbein and Kahn (2018)). Lindley and Machin (2016) nuance these views by showing that if the overall college wage premium stagnates, the postgraduate wage premium increases, suggesting SBTC now favors post-graduates. Deming (2017) also shows the rising returns to social instead of cognitive skills, and Deming and Noray (2020) focus on STEM works and exhibit the dampening of earning premium over time due to technology obsolescence. In a dynamic perspective, comparison of workers' wage evolution over their lifetime by cohorts has evidenced a flattening of the life-cycle profile for recent cohorts (Manovskii and Kambourov (2005), Guvenen et al. (2017) Kong et al. (2018)), which is particularly strong for college graduates (Rothstein (2020), Ashworth et al. (2021)). The reasons behind the flattening of lifetime income are debated. Jeong et al. (2015) point to the effect of demographic changes. Another possible explanation is the *scarring effect* brought about by the Great Recession that started at the end of the 2000s: a broad literature demonstrates the medium and long-term effects of exposure to bad economic conditions at the start of a career. Gregg and Tominey (2005) show youth exposure to unemployment carries a wage penalty twenty years later, Kahn (2010) focuses on college-educated workers and find persistent wage effects of graduating in a recession. In contrast, Brunner and Kuhn (2014) show blue-collar workers are penalized longer than white-collar workers because they have lower job mobility. Exploring the mechanisms behind the scarring effect, Liu et al. (2016) show skill mismatch between college graduates and firms increases during recessions, and explain career losses for 'unlucky cohorts'. Berge (2018) shows the initial mismatch fades thanks to job mobility, which lets unlucky cohorts catch up on luckier ones. Exploring a different mechanism, Kwon et al. (2010) show the intensity of worker promotions is correlated with the business cycle, which explains more than half of wage cohort effects.

The three strands of literature presented above all feed into the three chapters of this dissertation: chapter 1 empirically explores career mechanisms that result in different life-time income profiles by cohort, chapter 2 uses structural methods, namely matching models, that are particularly suited to understand supply and demand effects on the labor market at the aggregate level, to evaluate the simultaneous effect of an education expansion and technology change on wages, and chapter 3 develops a matching model to measure agents' expectations of individual returns to changes in their type. A summary of each of the dissertation chapters follows.

Chapter 1. The first chapter documents the flattening wage returns to experience for high education graduates in France between 1998 and 2017. I compare wage growth over the first

seven years in the labor market between three cohorts, who left school in 1998, 2004, and 2010. I document average wage growth by cohort and education level and find lower-educated workers (high school dropouts and high school graduates) experience flat wage profiles in the early years of their careers. However, High education graduates display the steepest profile for wage growth, but I find it is flattening between the 1998 and the 2010 cohort.

The French economy has undergone substantial changes over the period, as well as the French education system. As a result, the 1998, 2004, and 2010 cohorts enter the labor market under substantially different conditions: the 1998 cohort faces high (above 10%) unemployment, but counts relatively few high education graduates. The 2004 cohort enjoys low unemployment and high demand from firms but counts higher education graduates. Finally, the 2010 cohort enters the labor market amid the Great Recession, facing high unemployment and low growth. Because the French education expansion is still strong in the 2000s, encouraged by the creation of vocational bachelors and the implementation of the Bologna Process, the 2010 cohort counts substantially more high education graduates than its predecessors. This is likely to hurt their labor market prospects on several levels: first, they face reduced demand from firms. Second, high education graduates are more numerous than before. The impact of this second fact can be thought of from various perspectives: (Gaini et al. (2013); Dupray and Moullet (2010)). First, a degree could be a signal of individual productivity. An increase in the number of graduates then implies a decrease in their average individual quality, which can be reflected in a slower wage progression. A second approach considers the degree as a way of acquiring human capital. The diversification of the French higher education system, by modifying degrees' content, may have negatively impacted the acquisition of young graduates' human capital. Finally, a third approach based on the standard neoclassical model predicts a decline in the wages of young graduates if their number increases simply because the wage is equal to the marginal product of an additional worker. If firms exhibit decreasing returns to scale, every new worker lowers the average wage.

I set out to empirically study the reasons for differentiated wage growth in France since the end of the 1990s. Tod o so, I use the French 'Generations' surveys made available by the CEREQ (Centre d'Etudes et de Recherche sur les Qualifications). The surveys are presented as panel data and cover the working lives of school leavers in 1998, 2004, and 2010 for seven years, to provide a comprehensive overview of the integration of young people into the French labor market. I first decompose differences in average wage growth by occupation into an extensive and intensive margin. The extensive margin is driven by a composition effect resulting from differences in occupational shares between cohorts. The intensive margin rests on the change in annual wage growth by occupation. Occupations who display a

negative intensive margin are also the ones who exhibit a large and positive extensive margin. Indeed, occupations that know the most important slowdown in wage progression are also those which experience the greatest influx of graduates between 1998 and 2010. This is in line with a supply and demand interpretation of the wage growth slowdown, whereby an over-supply of new graduates prevents them from attaining their predecessors' wage levels. This interpretation suggests exploring mechanisms through which an increase in graduates' supply, along with stagnation in demand, impact early-career wage dynamics. I explore two mechanisms: promotion to managerial positions and degree-occupation mismatch. I show that obtaining a managerial position is accompanied by an increase in salary in the medium term. Hence, a decrease in the probability of obtaining such a position worsens the overall wage progression. This is consistent with findings by Kwon et al. (2010). I then examine the argument of Liu et al. (2016), who show that in the US, college graduates during the Great Recession suffered from a degraded degree-industry match, which led to persistently lower wage levels than their older peers. In France, I do not observe a worsening of mismatch (defined as the mean first-year wage level for a given degree major within a given occupation) between 1998 and 2010, but I find that its importance in determining future wages has increased between the 1998 and 2010 generations.

Chapter 2. Between the 1970s and today, many economies both in the developed and developing world have experienced an increase in their educated labor supply. As a result, the ratio of educated workers to uneducated workers present in labor markets has risen. This chapter seeks to understand the shift in labor supply education's impact on worker and firm matching and wages with a novel model of matching on the labor market. The model is structurally estimated on Portuguese matched employer-employee data. In doing so, I am able to quantify the impact of supply and demand changes on worker-firm allocation and wage structure.

To capture supply (from workers) and demand (from firms) mechanisms in the labor market, I build a static one-to-many matching model with transferable utility. Workers and firms differ with respect to their observed characteristics, which are summarized by a multi-dimensional type, as well as a stochastic shock that accounts for unobserved heterogeneity. A single firm matches with several workers, who constitute a bundle that forms its workforce. The surplus created by the match depends on the firms' observable characteristics as well as the workforce. The utility is transferable under the form of wages paid by the firm to the workers in its workforce. Firms seek to maximize total profit, which is additive in the difference of production and total wage bill, plus random shocks. Workers maximize their utility, which is additive in amenities, wage, and a random shock. Amenities embody

workers' inner preference for a given type of firm. At equilibrium, wages clear the market and each agent matches with their best option given wages. The model can generate a rich distribution of wages that depend both on workers' and firms' observable characteristics, as well as on the employed workforce. It also predicts equilibrium matching, which is the joint distribution of firms and workforces. Using both matching and wages, I can separately identify firm production from workers' amenities.

The framework offers more flexibility in estimation than classic supply and demand models developed in Katz and Murphy (1992) and Card and Lemieux (2001): it identifies worker preferences in addition to firm production, as well as varying production parameters over time. This is because by explicitly modeling firms' and workers' match choices, I can use both observed matching and observed wages, which brings more power to identification. The model is fitted to the data by assuming parametric forms for firm production and workers' amenities. I classify workers into two education levels, high school graduates and non-graduates, and three age groups, young, middle-aged, and senior. Firms are differentiated by their sector of activity. Following the literature, I choose a nested Constant Elasticity of Substitution (CES) function for production, with productivity parameters for each education level that vary between sectors. I assume worker preferences for firms depend on a worker's age, education level, and firm sector. Equipped with model predictions for matching and wages, I structurally estimate the model on matched employer-employee data. I estimate the model by maximum likelihood on the joint distribution of matching and wages, separately every three years.

The model developed in this chapter is related both to one-to-many assignment problems studied in mechanism design (Bikhchandani and Ostroy (2002), Vohra (2011)), and to one-to-one matching models used in family economics (Choo and Siow (2006)). This chapter bridges the gap between these two literatures: it extends one-sided assignments to two-sided matching, and generalizes one-to-one matching to one-to-many. Additionally, I extend the econometric framework of Choo and Siow (2006) and Galichon and Salanié (2021) to one-to-many matching.

I use the novel theoretical framework developed to study the Portuguese labor market between 1987 and 2017. I highlight three facts on the Portuguese labor market: first, the country operates a vast education expansion over the period, which translates in a dramatic increase in the relative supply of high school graduates relative to non-graduates on the labor market. Second, the high school wage premium decreases over the period. The high school wage premium is defined as the wage gap between workers who graduated from high school, and those who did not. The decrease in wage premium is particularly stark among young workers. Two opposite interpretations could be given for this fact. It could

be the consequence of a trade effect: Portugal joined the European Union in 1986, and because it had relatively more uneducated workers (workers who did not go to high school) in its labor force than other EU countries, a Heckscher-Ohlin model of trade would predict an increase in the export of goods that require uneducated labor to produce. Relative demand of uneducated to educated labor would rise, and the high school wage premium would decrease. The other interpretation rests on a supply effect: even if relative demand of educated to uneducated labor increases, the formidable education expansion occurring in Portugal in the 1990s and 2000s could take over and cause a decline in the high school wage premium. The model described above can distinguish which of the two, trade effect or education expansion effect, actually happened. Third, I measure worker-firm sorting, which is defined as the relative number of high school graduates over non-graduates in an age group employed in a given sector. The distribution of high school graduates versus non-graduates across industry sectors becomes highly unbalanced, in favor of services, and transports and communications, who employ an increasing share of high school graduates. The former two facts imply relative supply of high school graduates over non-graduates has grown faster than firms' relative demand for high school graduates over non-graduates. The latter suggests that sorting between workers and firms has evolved over the period: either because firms in services and transport and communications demand an increasing share of high school graduates, or because high school graduates' preference for these firms strengthens.

After estimating the model, I find that relative demand for high school graduates from firms in the Services, Manufacturing, and Transport & Communications sectors has increased dramatically over the period, starting in the early 2010s. This finding is in line with the skill-biased technological change hypothesis, rather than the trade hypothesis: it suggests an increase in the relative demand of educated to uneducated labor, that is outbalanced by the increase in relative supply. I also find that young and middle-aged high school graduates' preference for these industries has declined over time, while their share in production increases compared to senior workers. Compared to the classic supply and demand framework, the model offers two additional mechanisms whereby high school wages gaps stay large when a large number of high school-educated workers enter the labor market. First, a decrease in workers' amenities pressures wages upwards. Second, variation in young graduates' share in production compared to more senior high school graduates increases firm demands for the former compared to the latter. I perform several counterfactual exercises to assess the separate actions of changes in workers' demographics (both in education and age distribution), firm sector composition, firm demand through production parameters, and worker preferences, on sorting and wage premium. I find that changes in demographics are the main

positive drivers of changes in sorting. Changes in industry composition, firm demand, and worker preferences overall have a negative, if modest, effect on sorting. Wage premia by age group and industry are negatively affected by changes in worker demography and industry composition and positively affected by changes in firms' demand. These suggest changes in relative productivity in favor of high school graduates have driven the high school wage premium up, but cannot compensate for the large increase in the relative supply of graduates versus non-graduates.

Chapter 3. *Co-authored with Jeremy Fox and Alfred Galichon.* This chapter takes a different perspective from the two first chapters: instead of studying worker and firm matching in a static world, it asks how dynamic considerations influence matching. Indeed, on many matching markets, including the labor market but also the marriage market or startup ventures, agents account for the fact that their type may evolve in the future, either deterministically (for instance, workers age) or depending on whom they matched with (if a worker is employed in a given occupation, she will accumulate human capital in that occupation). To understand how these considerations influence partner or employer/employee choices, we introduce a tractable model of one-to-one dynamic matching. Agents have individual types, such as education and experience for workers, and industry and occupation for jobs. When deciding with whom to match, agents account for future expected returns that stem from a change in type. In turn, this change in type will affect returns from future matches. Each period a matching market takes place, where wages act as market-clearing prices. The goal of this chapter is to develop a useful off-the-shelf model of repeated matching games from the theory literature that generalizes static matching games to a dynamic setting. It also differs from the two previous chapters because the model developed is not only applicable to labor questions but could also apply in family economics or industrial organization. The repeated matching game with econometric errors can best be explained as the combination of two touchstone papers in the literature: Choo and Siow (2006) proposes an estimator for static matching games with logit errors. Rust (1987) proposes an estimator for single agent, dynamic discrete choice models, often using logit errors. What we do in this chapter is combine the two to obtain an estimator for dynamic matching games.

In our repeated matching framework, each agent has a state variable. Making a match, or remaining unmatched can affect the evolution of this agent state variable. Each period, agents participate in a matching market with prices or transfers for different matches. Given market-clearing prices, each agent selects the best partner in an, importantly, forward-looking manner. Next period the matching market reopens, new prices are stated and new matches form. A repeated matching game can have both individual and aggregate dynamics. At

the individual level, each agent is solving a single-agent dynamic programming problem, where each period the agent's action is to choose a partner to match with. At the aggregate level, the state variable of the matching market is the active agents' current set of types or state variables. This aggregate state variable evolves with the decisions of the individual agents and is summed up by a social planner Bellman equation. We first develop the model without econometric errors in surplus and then account for individual preferences under the form of econometric shock. In both cases, we explore two different methods to compute the aggregate equilibrium: one method rests on value function iteration on a grid to compute the social planner's value function, and the associated equilibrium on each point of the grid, and the second method uses a deep learning model to approximate a fixed point to the social planner's Bellman equation. One of our most important theoretical results is that a stationary equilibrium exists, both with and without econometric shocks: there is a mass of agent state variables such that, after optimal matches are chosen by forward-looking agents, the same masses of agent types occur. The existence of a stationary equilibrium does not depend on model parameters and lets the researcher optionally ignore aggregate dynamics by imposing that the matching game is at a stationary equilibrium. Focusing on the stationary equilibrium, we introduce yet two other methods to compute it: One method solves a system of nonlinear equations using a nonlinear programming solver. The second method reformulates the problem of finding a stationary equilibrium as a min-max problem and uses the Chambolle- Pock primal-dual algorithm to solve it. We show that both these methods can scale to problems with many agent types. In addition to computing a stationary equilibrium, we can extend the same estimators to structurally estimate parameters in the production of a match with an appropriate dataset. We then estimate geographic mobility costs for Swedish engineers in the 1970, and find that mobility imposes a sizeable penalty on production surplus.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.

Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *The Quarterly Journal of Economics*, 113(4):1055–1089.

Acemoglu, D. (2003). Patterns of Skill Premia. *The Review of Economic Studies*, 70(2):199–230.

Albrecht, J., Gautier, P. A., and Vroman, S. (2006). Equilibrium Directed Search with Multiple Applications. *The Review of Economic Studies*, 73(4):869–891.

Ashlagi, I. and Roth, A. E. (2012). New Challenges in Multihospital Kidney Exchange. *American Economic Review*, 102(3):354–359.

Ashworth, J., Hotz, V. J., Maurel, A., and Ransom, T. (2021). Changes across Cohorts in Wage Returns to Schooling and Early Work Experiences. *Journal of Labor Economics*, 39(4):931–964.

Autor, D., Goldin, C., and Katz, L. F. (2020). Extending the Race between Education and Technology. *AEA Papers and Proceedings*, 110:347–351.

Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the "other 99 percent". *Science*, 344(6186):843–851.

Autor, D. H. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing Inequality: Have Computers Changed the Labor Market? *The Quarterly Journal of Economics*, 113(4):1169–1213.

Azevedo, E. M. and Hatfield, J. W. (2018). Existence of Equilibrium in Large Matching Markets With Complementarities. SSRN Scholarly Paper ID 3268884.

Bagger, J. and Lentz, R. (2019). An Empirical Model of Wage Dispersion with Sorting. *The Review of Economic Studies*, 86(1):153–190.

Beaudry, P., Green, D. A., and Sand, B. M. (2015). The Great Reversal in the Demand for Skill and Cognitive Tasks. *Journal of Labor Economics*, 34(S1):S199–S247.

Becker, G. S. (1973). A Theory of Marriage: Part I. *Journal of Political Economy*, 81(4):813–846.

Berge, W. v. d. (2018). Bad start, bad match? The early career effects of graduating in a recession for vocational and academic graduates. *Labour Economics*, 53:75–96.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.

Bikhchandani, S. and Ostroy, J. M. (2002). The Package Assignment Model. *Journal of Economic Theory*, 107(2):377–406.

Blair, P. Q. and Deming, D. J. (2020). Structural Increases in Demand for Skill after the Great Recession. *AEA Papers and Proceedings*, 110:362–365.

Bloom, N., Draca, M., and Van Reenen, J. (2016). Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity. *The Review of Economic Studies*, 83(1):87–117.

Blundell, R., Green, D. A., and Jin, W. (2022). The U.K. as a Technological Follower: Higher Education Expansion and the College Wage Premium. *The Review of Economic Studies*, 89(1):142–180.

Boerma, J., Tsyvinski, A., and Zimin, A. P. (2021). Sorting with Team Formation. *NBER Working Paper 29290*.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A Distributional Framework for Matched Employer Employee Data. *Econometrica*, 87(3):699–739.

Bonnet, O., Galichon, A., Hsieh, Y.-W., O'Hara, K., and Shum, M. (2015). Yogurts Choose Consumers? Estimation of Random Utility Models via Two-Sided Matching. SSRN Scholarly Paper ID 2928876.

Bound, J. and Johnson, G. (1992). Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations. *The American Economic Review*, 82(3):371–392.

Brunner, B. and Kuhn, A. (2014). The impact of labor market entry conditions on initial job assignment and wages. *Journal of Population Economics*, 27(3):705–738.

Burdett, K. and Mortensen, D. T. (1998). Wage Differentials, Employer Size, and Unemployment. *International Economic Review*, 39(2):257–273.

Burdett, K., Shi, S., and Wright, R. (2001). Pricing and Matching with Frictions. *Journal of Political Economy*, 109(5):1060–1085.

Bárány, Z. L. and Siegel, C. (2018). Job Polarization and Structural Change. *American Economic Journal: Macroeconomics*, 10(1):57–89.

Cahuc, P., Postel-Vinay, F., and Robin, J.-M. (2006). Wage Bargaining with On-the-Job Search: Theory and Evidence. *Econometrica*, 74(2):323–364.

Card, D. and DiNardo, J. E. (2002). Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles. *Journal of Labor Economics*, 20(4):733–783.

Card, D., Heining, J., and Kline, P. (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics*, 128(3):967–1015.

Card, D. and Lemieux, T. (2001). Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis. *The Quarterly Journal of Economics*, 116(2):705–746.

Che, Y.-K., Kim, J., and Kojima, F. (2019). Stable Matching in Large Economies. *Econometrica*, 87(1):65–110.

Chiappori, P.-A., Costa-Dias, M., Crossman, S., and Meghir, C. (2020a). Changes in Assortative Matching and Inequality in Income: Evidence for the UK. *Fiscal Studies*, 41(1):39–63.

Chiappori, P.-A., Dias, M. C., and Meghir, C. (2020b). Changes in Assortative Matching: Theory and Evidence for the US. NBER Working Paper 26932.

Chiappori, P.-A., McCann, R. J., and Nesheim, L. P. (2010). Hedonic Price Equilibria, Stable Matching, and Optimal Transport: Equivalence, Topology, and Uniqueness. *Economic Theory*, 42(2):317–354.

Chiappori, P.-A., Oreffice, S., and Quintana-Domeque, C. (2012). Fatter Attraction: Anthropometric and Socioeconomic Matching on the Marriage Market. *Journal of Political Economy*, 120(4):659–695.

Chiappori, P.-A., Salanié, B., and Weiss, Y. (2017). Partner Choice, Investment in Children, and the Marital College Premium. *American Economic Review*, 107(8):2109–2167.

Choné, P. and Kramarz, F. (2021). Matching Workers' Skills and Firms' Technologies: From Bundling to Unbundling. *CREST Working Paper*.

Choo, E. and Siow, A. (2006). Who Marries Whom and Why. *Journal of Political Economy*, 114(1):175–201.

Ciscato, E., Galichon, A., and Goussé, M. (2020). Like Attract Like? A Structural Comparison of Homogamy across Same-Sex and Different-Sex Households. *Journal of Political Economy*, 128(2):740–781.

Deming, D. J. (2017). The Growing Importance of Social Skills in the Labor Market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Deming, D. J. and Noray, K. (2020). Earnings Dynamics, Changing Job Skills, and STEM Careers. *The Quarterly Journal of Economics*, 135(4):1965–2005.

Diamond, P. A. (1982). Wage Determination and Efficiency in Search Equilibrium. *The Review of Economic Studies*, 49(2):217–227.

Doepke, M. and Gaetani, R. (2020). Why Didn't the College Premium Rise Everywhere? Employment Protection and On-the-Job Investment in Skills. *IZA Working Paper*.

Dupray, A. and Moullet, S. (2010). Offre éducative, valorisation des diplômes et effets de composition : deux générations de sortants de l'université au tournant des années 2000. *Revue d'economie politique*, Vol. 120(5):845–880.

Dupuy, A. and Galichon, A. (2014). Personality Traits and the Marriage Market. *Journal of Political Economy*, 122(6):1271–1319.

Dupuy, A. and Galichon, A. (2022). A note on the estimation of job amenities and labor productivity. *Quantitative Economics*, 13(1):153–177.

Dupuy, A., Galichon, A., Jaffe, S., and Kominers, S. D. (2020). Taxation in Matching Markets. *International Economic Review*, 61(4):1591–1634.

Dupuy, A., Kennes, J., and Sun Lyng, R. (2021). Job Amenities in the Market for CEOs. *Working Paper*.

Eeckhout, J. and Kircher, P. (2011). Identifying Sorting–In Theory. *Review of Economic Studies*, 78(3):872–906.

Ekeland, I., Heckman, J., and Nesheim, L. (2004). Identification and Estimation of Hedonic Models. *Journal of Political Economy*, 112(S1):S60–S109. Publisher: The University of Chicago Press.

Fox, J., Yang, C., and Hsu, D. H. (2018). Unobserved Heterogeneity in Matching Games. *Journal of Political Economy*, Vol 126, No 4.

Fox, J. T. (2010). Identification in matching games. *Quantitative Economics*, 1(2):203–254.

Fox, J. T. and Bajari, P. (2013). Measuring the Efficiency of an FCC Spectrum Auction. *American Economic Journal: Microeconomics*, 5(1):100–146.

Gaini, M., Leduc, A., and Vicard, A. (2013). Peut on parler de générations sacrifiées ? Entrer sur le marché du travail dans une période de mauvaise conjoncture économique. *Économie et Statistique*, n 462-463.

Gale, D. and Shapley, L. S. (1962). College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1):9–15.

Galenianos, M. and Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, 144(2):445–471.

Galichon, A. (2016). *Optimal Transport Methods in Economics*. Princeton University Press.

Galichon, A. (2021). The unreasonable effectiveness of optimal transport in economics. *arXiv:2107.04700 [econ, q-fin]*.

Galichon, A. and Hsieh, Y.-W. (2018). Aggregate Stable Matching with Money Burning. SSRN Scholarly Paper ID 2887732.

Galichon, A. and Salanié, B. (2021). Cupid's Invisible Hand: Social Surplus and Identification in Matching Models. *The Review of Economic Studies*.

Garibaldi, P., Moen, E. R., and Sommervoll, D. E. (2016). Competitive on-the-job search. *Review of Economic Dynamics*, 19:88–107.

Goldin, C. and Katz, L. (2008). *The Race Between Education and Technology.* Belknap Press for Harvard University Press.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, 104(8):2509–2526.

Gregg, P. and Tominey, E. (2005). The wage scar from male youth unemployment. *Labour Economics*, 12(4):487–509.

Gretsky, N., Ostroy, J., and Zame, W. (1992). The Nonatomic Assignment Model. *Economic Theory*, 2:103–27.

Gualdani, C. and Sinha, S. (2022). Partial Identification in Matching Models for the Marriage Market. *arXiv:1902.05610 [econ]*.

Guvenen, F., Kaplan, G., Song, J., and Weidner, J. (2017). Lifetime Incomes in the United States over Six Decades. NBER Working Paper 23371.

Hagedorn, M., Law, T. H., and Manovskii, I. (2017). Identifying Equilibrium Models of Labor Market Sorting. *Econometrica*, 85(1):29–65.

Hatfield, J. W. and Milgrom, P. R. (2005). Matching with Contracts. *American Economic Review*, 95(4):913–935.

Hershbein, B. and Kahn, L. B. (2018). Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings. *American Economic Review*, 108(7):1737–1772.

Hicks, J. R. (1932). The Theory of Wages. *London Macmillan*.

Hosios, A. J. (1990). On the Efficiency of Matching and Related Models of Search and Unemployment. *The Review of Economic Studies*, 57(2):279–298.

Jeong, H., Kim, Y., and Manovskii, I. (2015). The Price of Experience. *American Economic Review*, 105(2):784–815.

Julien, B., Kennes, J., and King, I. (2000). Bidding for Labor. *Review of Economic Dynamics*, 3(4):619–649.

Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2):303–316.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Kelso, A. S. and Crawford, V. P. (1982). Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica*, 50(6):1483–1504.

Kircher, P. (2009). Efficiency of Simultaneous Search. *Journal of Political Economy*, 117(5):861–913.

Kong, Y. C., Ravikumar, B., and Vandenbroucke, G. (2018). Explaining cross-cohort differences in life-cycle earnings. *European Economic Review*, 107:157–184.

Krueger, A. B. (1993). How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989. *The Quarterly Journal of Economics*, 108(1):33–60.

Krueger, D., Perri, F., Pistaferri, L., and Violante, G. L. (2010). Cross-sectional facts for macroeconomists. *Review of Economic Dynamics*, 13(1):1–14.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, 68(5):1029–1053.

Kwon, I., Milgrom, E. M., and Hwang, S. (2010). Cohort Effects in Promotions and Wages: Evidence from Sweden and the United States. *The Journal of Human Resources*, 45(3):772–808.

Lentz, R., Piyapromdee, S., and Robin, J.-M. (2018). The Anatomy of Sorting: Evidence from Danish Register Data. Working Paper.

Lindenlaub, I. (2017). Sorting Multidimensional Types: Theory and Application. *The Review of Economic Studies*, 84(2):718–789.

Lindley, J. and Machin, S. (2016). The Rising Postgraduate Wage Premium. *Economica*, 83(330):281–306.

Lise, J., Meghir, C., and Robin, J.-M. (2016). Matching, sorting and wages. *Review of Economic Dynamics*, 19:63–87.

Lise, J. and Postel-Vinay, F. (2020). Multidimensional Skills, Sorting, and Human Capital Accumulation. *American Economic Review*, 110(8):2328–2376.

Lise, J. and Robin, J.-M. (2017). The Macrodynamics of Sorting between Workers and Firms. *American Economic Review*, 107(4):1104–1135.

Liu, K., Salvanes, K. G., and Sørensen, E. O. (2016). Good skills in bad times: Cyclical skill mismatch and the long-term effects of graduating in a recession. *European Economic Review*, 84:3–17.

Lopes de Melo, R. (2018). Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence. *Journal of Political Economy*, 126(1):313–346.

Manning, A. (2021). Monopsony in Labor Markets: A Review. *Industrial and Labor Relations Review*, 74(1).

Manovskii, I. and Kambourov, G. (2005). Accounting for the Changing Life-Cycle Profile of Earnings. 2005 Meeting Paper 231, Society for Economic Dynamics.

McCall, J. J. (1970). Economics of Information and Job Search. *The Quarterly Journal of Economics*, 84(1):113–126.

Menzel, K. (2015). Large Matching Markets as Two-sided Demand Systems. *Econometrica*, 83(3):897–941.

Menzio, G. and Moen, E. R. (2010). Worker replacement. *Journal of Monetary Economics*, 57(6):623–636. Publisher: Elsevier.

Menzio, G. and Shi, S. (2011). Efficient Search on the Job and the Business Cycle. *Journal of Political Economy*, 119(3):468–510. Publisher: The University of Chicago Press.

Mortensen, D. (1970). Job Search, the Duration of Unemployment, and the Phillips Curve. *American Economic Review*, 60(5):847–62.

Mortensen, D. T. (1982). The Matching Process as a Noncooperative Bargaining Game. In *The Economics of Information and Uncertainty*, pages 233–258. University of Chicago Press.

Mortensen, D. T. (2003). *Wage Dispersion: Why Are Similar Workers Paid Differently?* Zeuthen Lectures. MIT Press, Cambridge, MA, USA.

Moscarini, G. and Postel-Vinay, F. (2018). The Cyclical Job Ladder. *Annual Review of Economics*, 10(1):165–188.

Peyré, G. and Cuturi, M. (2020). Computational Optimal Transport. *arXiv:1803.00567 [stat]*. arXiv: 1803.00567.

Pissarides, C. A. (1985). Dynamics of Unemployment, Vacancies and Real Wages with Trade Unions. *The Scandinavian Journal of Economics*, 87(2):386–403.

Postel–Vinay, F. and Robin, J.-M. (2002). Equilibrium Wage Dispersion with Worker and Employer Heterogeneity. *Econometrica*, 70(6):2295–2350.

Rani, U. and Furrer, M. (2016). Decomposing income inequality into factor income components: Evidence from selected G20 countries. ILO Research Paper No. 15, International Labor Office.

Roth, A. E. (1984). The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory. *Journal of Political Economy*, 92(6):991–1016.

Rothstein, J. (2020). The Lost Generation? Labor Market Outcomes for Post Great Recession Entrants. NBER Working Paper 27516.

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, 55(5):999–1033.

Schaal, E. (2017). Uncertainty and Unemployment. *Econometrica*, 85(6):1675–1721.

Shapley, L. S. and Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130.

Shimer, R. and Smith, L. (2000). Assortative Matching and Search. *Econometrica*, 68(2):343–369.

Sinha, S. (2018). Identification in One-to-One Matching Models with Nonparametric Unobservables. *TSE Working Paper*.

Song, J., Price, D. J., Guvenen, F., Bloom, N., and von Wachter, T. (2019). Firming Up Inequality. *The Quarterly Journal of Economics*, 134(1):1–50.

Tsuyuhara, K. (2016). Dynamic Contracts with Worker Mobility Via Directed on-the-Job Search. *International Economic Review*, 57(4):1405–1424.

Valletta, R. G. (2016). Recent Flattening in the Higher Education Wage Premium: Polarization, Skill Downgrading, or Both? NBER Working Paper 22935.

Verdugo, G. (2014). The great compression of the French wage structure, 1969–2008. *Labour Economics*, 28:131–144.

Villani, C. (2009). *Optimal Transport: Old and New*. Springer.

Vohra, R. V. (2011). *Mechanism Design: A Linear Programming Approach*. Econometric Society Monographs. Cambridge University Press, Cambridge.

Woodcock, S. D. (2008). Wage differentials in the presence of unobserved worker, firm, and match heterogeneity. *Labour Economics*, 15(4):771–793.

CHAPTER 1

The Flattening Returns to Experience for Higher Education Graduates in France

An Occupational Analysis

Abstract

This chapter documents the falling wage returns to experience between cohorts entering the French labor market in 1998 and 2010, as experienced by higher education graduates in their early careers. Returns remain stable for high school dropouts and high school graduates. I decompose differences in average wage growth by occupation into an extensive and intensive margin. The extensive margin is driven by a composition effect resulting from differences in occupational shares between cohorts. The intensive margin rests on the change in annual wage growth by occupation. Occupations that display a negative intensive margin are also the ones that exhibit a large and positive extensive margin. This finding is consistent with decreasing returns to each new graduate in a given occupation. I then study two mechanisms behind the wage growth slow down: access to managerial positions and impact of initial match quality. I find access to managerial positions is more infrequent for recent cohorts. I also find that initial match quality has not worsened between the 1998 and 2010 cohorts, but its impact on future wages has become more enduring.

1 Introduction

Wage growth on workers' career paths, or returns to experience, have been the subject of a substantial economic literature, especially since the recession that affected the world economy in the late 2000s and early 2010s. A subset of this literature, to which this chapter belongs, focuses on individuals' early careers and compares returns to experience by cohort, or generation, i.e. a set of individuals who entered the labor market at the same time. In this chapter, I compare the wage growth between three cohorts, who entered the French labor market in 1998, 2004, and 2010, over their first seven years of career. I document average wage growth by education level and find lower-educated workers (high school dropouts and high school graduates) experience flat wage profiles in the early years of their careers. High education graduates display the steepest profile for wage growth, but I find it is flattening between the 1998 and the 2010 cohorts.

Over the period studied in this article, from the end of the 1990s to the second half of the 2010s, the French economy has undergone two recessions at the beginning and end of the 2000s: the high unemployment rate at the beginning of the period (above 10% of the total active population) fell sharply during the 2000s before rising again from 2008 onwards (without, however, reaching its previous level). GDP grew steadily throughout the period, except in 2008 and 2009. At the same time, the French education system evolved in the early 2000s, as vocational bachelors were created and the Bologna process was implemented. The latter reform reorganized the French higher education system into three levels: Bachelor (3 years), Masters (2 years), and Doctorate (3 years). Both reforms contributed to multiplying the number of university graduates entering the labor market and diversified their profiles.

Two mechanisms are at work between 1998 when the first cohort studied enters the labor market, and 2017, the last year the third cohort is observed: demand, reflected in job opportunities that vary according to firms' needs, and supply, since the population's educational composition changes between 1998 and 2017, with an increase in the share of higher educated workers. On the supply side, the impact of the increase in higher education graduates on wage levels can be thought of from various perspectives (Gaini et al. (2013); Dupray and Moullet (2010)). First, a degree could be a signal of individual quality which is unobserved by the analyst but observed by the employer, who adjusts the wage accordingly. If we assume that the unobserved quality is distributed in the same way among each cohort, an increase in the number of graduates implies a decrease in their average unobserved quality, which can be reflected in a slower wage progression. In the US, the decrease in graduates' quality due to an

education expansion is evidenced in Carneiro and Lee (2011). A second approach considers the degree as a means of acquiring human capital: in general, high human capital would imply high wages, because human capital makes workers more productive. If the level of human capital conferred by the degree does not change, wages should not change, regardless of the number of graduates. But the diversification of French higher education, by modifying degrees' content, may have negatively impacted the acquisition of human capital of young graduates. Finally, a third approach based on the standard neoclassical model and developed by Katz and Murphy (1992) and Card and Lemieux (2001) predicts a decline in the wages of young graduates if their number increases without any change in demand, even if neither the unobserved quality nor the content of the degree changes. This approach postulates that the wage of an employee is simply on his or her marginal product. If firms produce with decreasing returns to scale, an additional employee has a smaller marginal product than the employees that were hired before she was. In a context of a strong increase in the number of graduates on the labor market without a comparable increase in demand for the most highly educated, the diminishing marginal returns approach therefore anticipates lower wages for the latest arrivals, i.e. young graduates. This last approach could also impact the demand side: if firms' demand for young high educated workers drops, it negatively affects their wage.

This chapter sets out to empirically study the reasons for differentiated wage growth in France since the end of the 1990s. I first document early career wage progression in France between 1998 and 2017 and show that it changes differently depending on education level and occupations. Indeed, occupations that know the most important slowdown in wage progression are also those which experience the greatest influx of graduates between 1998 and 2010. This is in line with a supply and demand interpretation of the wage growth slowdown, whereby an over-supply of new graduates prevents them from attaining their predecessors' wage levels. This interpretation suggests exploring mechanisms through which an increase in graduates' supply, along with stagnation in demand, impact early-career wage dynamics. I explore two mechanisms: promotion to managerial positions and degree-occupation mismatch. I show that obtaining a managerial position is accompanied by an increase in salary in the medium term. Hence, a decrease in the probability of obtaining such a position worsens the overall wage progression. This is consistent with findings by Kwon et al. (2010). I then examine the argument of Liu et al. (2016), who show that in the US, college graduates during the Great Recession suffered from a degraded degree-industry match, which led to persistently lower wage levels than their older peers. In France, I do not observe a worsening of mismatch (defined as the mean first-year wage level for a given degree major within a given occupation) between 1998 and 2010, but I find that its importance in determining

future wages has increased between the 1998 and 2010 generations.

I use the French 'Generations' surveys made available by the CEREQ (Centre d'Etudes et de Recherche sur les Qualifications). The surveys are presented as panel data and cover the working lives of school leavers in 1998, 2004, and 2010 for seven years, to provide a comprehensive overview of the integration of young people into the French labor market. The surveys show that the 2010 cohort (defined by individuals who left the education system in 2010, regardless of their age) has experienced a more difficult situation than its predecessors: three years after they entered working life, their unemployment rate was 22%, compared to 11% for the 1998 generation in 2001 (Epiphane et al. (2019)). The median wage in the first year on the labor market is higher for the 2010 cohort than for the 1998 cohort: 1265 versus 1090 in constant euros, base 2015. However, the 2010 cohort experiences a slower wage growth than the 1998 cohort: after seven years median wages are 1510 and 1500, respectively. Besides, the median wage after seven years is higher for the 1998 cohort than the 2010 cohort for higher education graduates, indicating a strong slowdown in salary progression for the highly educated.

Related literature. The present analysis relates to several literatures. First, there exists a wide literature on wage inequality by education level, that usually attributes the increasing education wage premium of the 1070s and 1980s to skill-biased technological change (Bound and Johnson (1992), Katz and Murphy (1992), Card and Lemieux (2001)). The subsequent flattening of the education wage premium in the 1990s and 2000s (Card and DiNardo (2002)) is either explained through education expansions experiences in Europe (see Verdugo (2014) on France and Blundell et al. (2022) on the UK), on employment protections that dampen skill-biased technological change (Doepke and Gaetani (2020) in Germany), or to a simple reversal in firms demand for cognitive skill (Beaudry et al. (2015), Valletta (2016)). On the opposite side of the spectrum Blair and Deming (2020) and Hershbein and Kahn (2018) argue the Great Recession has accelerated skill-biased technological change rather than slown it down. Finally Lindley and McIntosh (2015) and Lindley and Machin (2016) evidence that if the wage structure has compressed overall, wage inequality within college graduates has increased depending on graduates' chosen major and postgraduates studies. Second, several papers have found evidence of a flattening life-cycle wage profile: Manovskii and Kambourov (2005) document deflating life-cycle earnings for men in the US since the 1960s, which Jeong et al. (2015) explain entirely through the demographic changes (i.e. the worker supply side) that occurred over the period. Similarly, Guvenen et al. (2017) document stagflation of men's lifetime income in the US and Kong et al. (2018) show that the labor

earnings 'multiplier' between 25 and 55 has decreased from 4 to 2.6 between 1940 and 1980 for college-educated workers in the US. To explain this fact, they build a model of schooling and human capital accumulation on the job and find that the decrease is entirely due to the increasing price of skill: because high skilled workers are more demanded by firms, college enrolment has increased, which drove up levels of human capital at career start, resulting in slower earning growth afterward. Over the shorter term in France, I find no evidence of this mechanism: average starting wages for high education graduates in 1998, 2004, and 2010 are the same, which indicates younger cohorts do not have a human capital advantage over older ones. In related work, Ashworth et al. (2021) evidence the decrease in men's wage growth between 1979 and 1997 in the US, especially for college graduates. Their model of schooling and work decision points to composition effects both in terms of observable and unobservable skills to explain variation across cohorts. Third, this chapter relates to the literature on labor market entry conditions and their impact on workers' career trajectories and medium and long-term earnings (see von Wachter (2020) for a complete survey): Gregg and Tominey (2005) finds that exposure to unemployment still carries a wage penalty for workers after 20 years. Stevens (2008) attenuate this finding by showing that labor market outcomes of low and medium-skilled workers are not very vulnerable to economic conditions at the start of the career. Focusing on college graduates in the US Oreopoulos et al. (2012) show negative and persistent effects graduating in a recession on wage. They also show that workers partially recover through mobility towards better-paying firms. This finding is related to Liu et al. (2016)'s work on mismatch in the US: they show that in a recession mismatch between college majors and firms' sectors increases, which drives wages down and creates a persistence in the decrease. Job mobility allows to dampen the effect over time through new matches with better firms(Fredriksson et al. (2018), Berge (2018)). In this chapter, I find on the contrary that in France mismatch does not worsen for younger cohorts, but its weight on future wage determination intensifies. Job mobility is not sufficient to alleviate this intensification. Brunner and Kuhn (2014) differentiate their analysis by socio-professional category in Austria and show that the recession particularly affected blue-collar workers, as they are stuck in low-quality jobs longer than white-collar workers. I also take an occupational angle in this chapter, but do not find white-collar occupations to be less affected than blue-collar occupations: rather, I find that the occupations that suffer the most from deteriorated wage growth are the ones that experience the largest influx of new graduates. Not all work finds long and persistent effects on entering the labor market in a recession: for instance Berge and Brouwers (2017) shows the wage penalty lasts about fours years for high education graduates in the Netherlands. Finally, this analysis is closely related to Rothstein (2020) who shows wages decline in the short and medium-term in the

US, starting for cohorts who enter the labor market in 2005: I find a similar trend in France since the middle cohort I observe, who enters the labor market in 2004, already displays deteriorated wage growth for high education graduates compared to the 1998 cohort, even if it enters the labor market in times of good economic conditions.

Section 2 describes the economic context in which the 1998, 2004, and 2010 cohorts entered, as well as the data from the Generations Surveys and the main variables of interest. I also introduce the empirical framework. Section 3 presents a decomposition of wage growth by socio-professional category. In section 4, I present two mechanisms of the slowdown in wage growth among the most highly educated. Finally, section 5 presents robustness tests and section 6 concludes.

2 Data and Empirical Strategy

2.1 The French labor market between 1998 and 2017

Wage changes examined in this article are part of general trends on the French labour market between 1998 and 2017. Table 1.1 uses INSEE census data to provide a general overview of changes in the composition of the educational levels and occupations of the working population between 1999 and 2011. On the supply side, the share of higher education graduates in 1999 was 24.6% of the working age population. In 2011, this share is 36.4%, a gain of almost 3 million individuals. The share of individuals with a high school degree has also increased, but to a lesser extent. The evolution of demand for each education levels is more difficult to assess, and I approximate it by the share of each occupation in the general population. The occupations whose numbers increased the most between 1998 and 2011 are managers and higher intellectual occupations (MHIO) and intermediate occupations (IO). Employees and craftsmen, shopkeepers and business owners saw their numbers stagnate, while they decreased for farmers and plant workers. In 1999, the MHIO positions were mainly occupied by higher education graduates (76.3%), and this share increased in 2011 (82%). However, the share of higher education graduates in intermediate occupations has also increased, from 43.7% to 55.1%. In absolute terms, this increase even surpasses that of the MHIOs: 1,026 thousand individuals compared to 920. The strong link between tertiary graduates and MHIOs in 1999 has thus been weakened in favour of IOs in 2011. Three mechanisms may jointly explain this evolution: firstly, the nature of the tasks required within the MHIOs and IOs may have changed. The literature on job polarisation in France (Albertini et al. 2017; Patel 2020) associates MHIOs with abstract tasks, and IOs

with routine tasks, with higher education graduates being most suited to abstract tasks. If it is the case that task content required in IOs positions tends towards more abstraction, the demand for higher education graduates should increase within this occupation. Secondly, it may be that supply of higher education graduates in 2011 is above demand from MHIOs, pushing them towards intermediate occupations by default. Finally, either the content of higher education degrees, or the graduates themselves may have changed between 1999 and 2011, and higher education graduates who entered the labour market between 1999 and 2011 are more productive at performing routine than abstract tasks. Understanding how the type of tasks associated with different occupations has evolved over time is already a focus in the polarization literature, hence I choose to explore the last two mechanisms in this paper.

Table 1.1: Education Levels by Occupations within French active population in 1999 and 2011

	1999			2011			Difference	
Occupation	Nb (k)	% HS	% HE	Nb (k)	% HS	% HE	HS (k)	HE (k)
Farmers	532	16.1	7.5	344	28.3	18.4	12	23
Craftmen, retailers, business owners	1 407	15.0	14.9	1 367	21.4	23.9	82	116
Top managers, highly qualified professionals	2 802	10.6	76.3	3 726	9.2	82.0	47	920
Mid-level managers	5 100	21.8	43.7	5 905	21.1	55.1	134	1 026
Employees	6 587	16.7	10.5	6 522	24.8	20.5	516	646
Factory workers	5 827	6.2	2.7	5 162	15.1	6.8	418	193
Total	22 255	14.2	24.6	23 026	19.0	36.4	1 210	2 925

Nb (k): Number of individuals in thousands

HS: High School degree, HE: Higher Education degree

Lastly, changes in unemployment rate also reveals disparity between supply and demand. Public data from INSEE (Institut National de la Statistique et des Etudes Economiques) on unemployment rate by level of education show a systematically higher rate for high school graduates than for higher education graduates. Moreover, the unemployment rate increased between 1997 and 2017 for non-graduates, from 14% to around 17%, and high school graduates (around 12%), while it decreased for higher education graduates (from 7.5% to around 4.5%). This observation is consistent with an adjustment by unemployment among

the less educated during a fall in demand, as the minimum wage prevents any adjustment of wages (Gaini, Leduc, and Vicard 2013). For higher education graduates, on the other hand, any friction between supply and demand is rather reflected in wage levels rather than in the unemployment rate.

2.2 The data

The Generations Surveys are presented in the form of a panel: each observation corresponds to the activity of an individual (employment or unemployment) over a given period, called a 'sequence' (or spell). CEREQ conducts its surveys on a given cohort every two or three years. For instance, the 2010 cohort is surveyed in 2013, 2015 and 2017. Only individuals who responded to all three surveys are considered here. The three surveys are unequal in terms of the number of individuals surveyed: there are twice as many individuals surveyed from 1998 versus the 2010 cohort. To account for these differences, and any selection effect that may arise from attrition, the Generation Surveys provide the analyst with a weighting per individual so that each survey is representative of the population of young French workers. I adapt this weighting in two ways: first, I normalise it so that each of the generations 1998, 2004 and 2010 has the same weight. Second, since the data are presented as individual-spell observations, individuals who change spell frequently is greatly increased in the analysis (interim workers for instance). To avoid this giving them too much weight, I weigh spells of individuals who change status several times a year according to spell length. The entire analysis will be weighted by these modified weights.

Table 1.2: Number of individuals and spells by cohort

	Gen 1998	Gen 2004	Gen 2010
Number of individuals	13 673	9 633	7 500
Number of spells	63 965	45 343	34 730
Number of employment spells	27 618	21 576	15 533

The analysis focuses on employment spells and starting (or entry) wage obtained by young workers hired at the beginning of these spells, the changes in which is compared between the 1998, 2004 and 2010 cohort by level of education. Each individual first spell starts the month after graduation or after they left school if they did not graduate. The surveys also provide the last wage obtained at the end of each spell, but no intermediary wage. I choose to focus on entry wage, because this is invariant to the duration of the employment spell. Using the INSEE consumer price index series, I compute wages in constant prices in

euro 2017.

I exclude from the analysis spells in which individuals are under 16 years old, as well as the employment spells for which the monthly starting wage is less than € 200 or more than € 20,000. The analysis focuses on job spell for which the location (at the 'department' level), firm sector and occupation are known. I consider only sequences in metropolitan France, between year 1 and year 8 of each cohort.

The main characteristics of the individuals are described in Table 1.3: there are no major differences between cohorts in terms of the average age just after leaving the education system, the gender distribution, or the average number of spells after seven years.

Table 1.3: Age, gender, and individual number of spells by cohorts

	Gen 1998	Gen 2004	Gen 2010
Average age at entry on labor market	21.6	21.2	21.3
% Men	0.51	0.53	0.51
Average number of spells	4.9	5.3	5.1
Average number of employment spells	2.1	2.4	2.2

I consider two main dimensions of individuals and their employment: educational attainment and occupation. I group individuals into four education levels: no degree (left the education system without finishing secondary school), secondary education (obtained either a general high school degree, or a vocational degree), short higher education (obtained a degree in less than four years, either a bachelor or a technical degree), and long higher education (obtained a degree in more than four years, either masters or PhD). Table 1.4 presents the composition of each cohort by level of education: the proportion of long higher education graduates (more than four years of higher education) is greater in the 2010 generation than in the 2004 and 1998 cohorts, while the proportions of short tertiary graduates (between one and three years of higher education) and secondary school graduates (CAP, BEP or Baccalauréat) are lower. The proportion of individuals without a diploma (having left school with a brevet level) is higher in the 2010 cohort. The Generations surveys therefore show a polarisation of educational provision between 1998 and 2010.

Table 1.4: Education level shares by cohort

Education level (%)	Gen 1998	Gen 2004	Gen 2010
No degree	8.9	7.9	17.2
High school degree	52.3	53	42.7
Short higher education degree	28.1	27.6	23.3
Long higher education degree	10.7	11.5	16.9
Total	100	100	100

Table 1.5 shows the decomposition of occupations for first job by cohort. The share of managers and upper occupations increased between the 1998 and 2010 cohorts. The share of intermediate occupations has also increased, at a faster pace. As in the general population, the share of blue-collar workers has decreased, and the share of white-collar workers has stagnated. Because farmers represent too small a share of the employment spell, these spells are excluded from the rest of the analysis.

Table 1.5: Occupation shares by cohort

Occupation (%)	Gen 1998	Gen 2004	Gen 2010
Farmers	0.7	0.2	-
Craftmen, retailers, business owners	1.1	0.7	0.1
Top managers, highly qualified professionals	12.1	11.4	18.6
Mid-level managers	25.2	29.9	30.3
Employees	28.2	26.9	26.4
Factory workers	32.7	30.9	24.6
Total	100	100	100

2.3 Strategy

Figure 1.1 shows the evolution of average starting wage by cohort and education level reported by individuals on the Generation Surveys. Education levels are grouped into four categories: no degree (individuals who left school at brevet level), secondary degree, short tertiary diploma (graduates of a higher education degree in three years or less), and long tertiary degree (graduates of a higher education degree in four or more years). This graph shows no significant difference in how starting wages growth over time between cohorts for three out of four levels of education: non-graduates, secondary school graduates and short tertiary graduates. On the other hand, long-term tertiary graduates' entry wage growth

differs between cohorts: the 2010 cohort experiences a slower growth than the 2004 and 1998 cohorts. This slowdown becomes more pronounced over time: the three cohorts begin their working lives with similar entry wage, and then diverge. While the 1998 cohort enjoys a significant average increase in starting wage from their second year on the labour market, the 2004 cohort's average starting wage only really increased after three years on the labour market, and that of the 2010 generation after four years. The result is that the 2010 generation is significantly behind its predecessors, a gap that persists beyond the 2010-2012 crisis period (years 1 to 3 for the 2010 generation), without any catching up taking place in the subsequent years available in the survey.

Figure 1.1: Average entry wage over time in euros, by cohort and education level

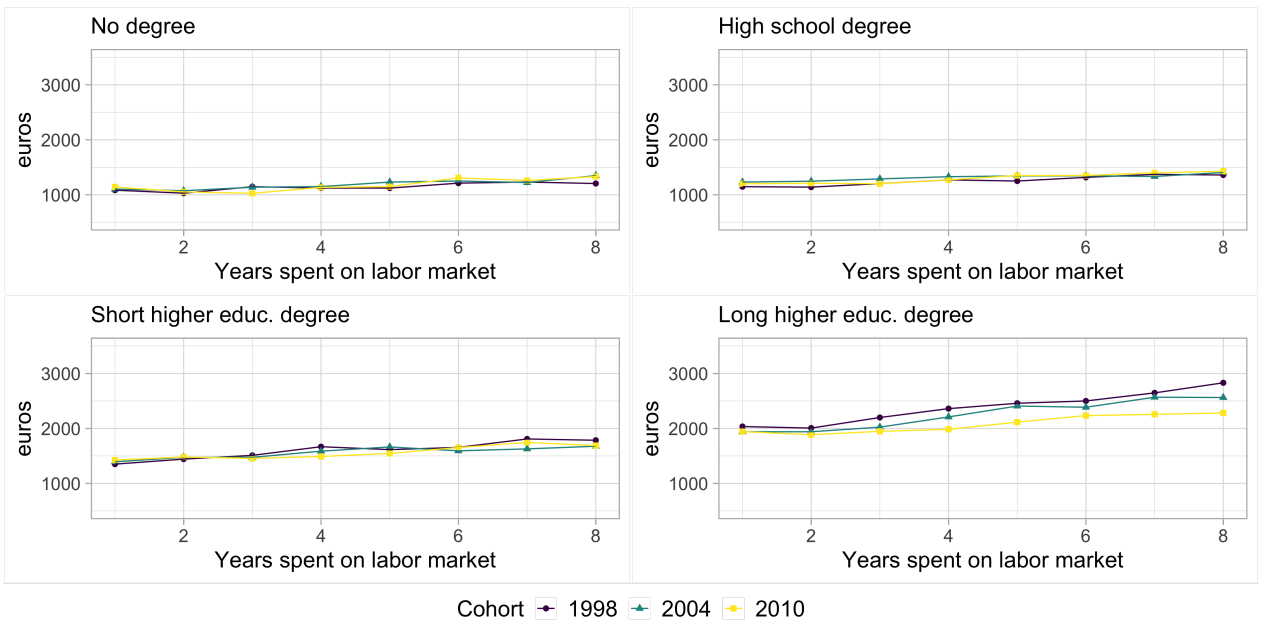

To understand the reasons behind the divergence observed in Figure 1.1, I use the following framework: individual i enters employment contract $j = J(i, t)$ in year t. Each contract j is characterised by the characteristics of the firm, such as the industry or region, but also by features specific to the individual's role in the company, such as occupation. The individual also displays specific characteristics such as cohort or level of education. Entry monthly wage under contract j in year t is w_{jt}. The regression below allows to decompose the evolution of average entry wages by cohorts and level of education, considering possible differences in the characteristics of contracts or individuals:

$$\log w_{ijt} = \sum_e \mathbb{1}_{[educ_i=e]}\beta_{eg} \times a_t + e_i + g_i + r_j + s_j + \epsilon_{ijt}, \tag{1.1}$$

where $\log w_{ijt}$ is the logarithm applied to entry wage, a_t is the number of years since leaving the education system (between 1 and 8), e_i is a fixed effect for education, g_i a fixed effect for the individual gender, r_j is a fixed effect for region and s_j for industry within which the contract takes place.

The estimator β_{eg} Is computed by education level e and cohort g. It measures the average increase in entry wages per year, for each cohort and education level, controlling for variations in gender, industry, and region. Comparison of estimators between cohorts is based on the following identification assumption: the distribution of unobserved heterogeneity is the same for all generations. This assumption will be maintained for the rest of the analysis.

One way of approaching the variations of β_{eg} between cohorts in a context of changing supply and demand is to decompose average entry wage growth not only by cohort and level of education, but also by occupation. In fact, by highlighting the heterogeneity of wage growth by occupations, I can identified two margins of divergence: an extensive margin and an intensive margin. The extensive margin highlights the variations in the share of new hires in each occupation, keeping wage evolution constant. The intensive margin focuses on variations in wage levels, holding constant the share of each occupation in new hires. This margin decomposition proceeds in two steps: first, define w^0_{ijt}, average entry wage cleaned fixed effects in the previous regression:

$$\log w^0_{ijt} = \sum_g \mathbb{1}_{[coh_i=g]} \sum_e \mathbb{1}_{[educ_i=e]} \hat{\beta}_{eg} \times a_t + \epsilon_{ijt}. \tag{1.2}$$

The second step is to project $\log w^0_{jt}$ onto the space of education and occupation by estimating the following regression by cohort:

$$\log w^0_{ijt} = \sum_e \mathbb{1}_{[educ_i=e]} \sum_g \mathbb{1}_{[occ_j=p]} \gamma_{egp} \times a_t + \epsilon_{ijt}. \tag{1.3}$$

The following decomposition is then carried out, for a given level of education e :

$$\hat{\beta}_{e,2010} - \hat{\beta}_{e,1998} = \sum_p n_{e,2010,p} \times \hat{\gamma}_{e,2010,p} - \sum_p n_{e,1998,p} \times \hat{\gamma}_{e,1998,p}, \tag{1.4}$$

where $n_{e,1998,p}$ et $n_{e,2010,p}$ are the respective proportions of each occupation p within the education level e and the 1998 and 2010 generations. Introducing the cross term $\sum_p n_{e,2010,p} \times \hat{\gamma}_{e,1998,p}$ we obtain:

$$\hat{\beta}_{e,2010} - \hat{\beta}_{e,1998} = \sum_p (n_{e,2010,p} - n_{e,1998,p}) \times \hat{\gamma}_{e,1998,p} \\ - \sum_p n_{e,1998,p} \times (\hat{\gamma}_{e,2010,p} - \hat{\gamma}_{e,1998,p}) \,. \tag{1.5}$$

The first term $(n_{e,2010,p} - n_{e,1998,p}) \times \hat{\gamma}_{e,1998,p}$ corresponds to the extensive margin: the share of the change in the slope of entry wage growth due to changes in share of occupation p within new hires. The second term $n_{e,2010,p} \times (\hat{\gamma}_{e,2010,p} - \hat{\gamma}_{e,1998,p})$ is an intensive margin: the share of the change in the slope strictly due to the change in the slope for specific occupation p, holding constant the share of each occupation in new hires. This decomposition seeks to separate a pure demand or composition effect (changes in the occupation of new hires between cohorts, i.e. the extensive margin) from a supply and demand equilibrium effect (changes in the distribution of education levels within individuals, captured by the intensive margin).

3 Results

Estimation results for equation (1.1) are presented in Table 1.6. Coefficients for entry-level wage growth are significant for all cohorts and levels of education. Hiring wages of individuals with no degree and high school graduates increase slightly during the first seven years on the labor market for all generations (about between 1.9% and 3.7% per year). Both short and long higher education graduates experience a more sustained growth in wage, but it is less pronounced for the 2004 and 2010 cohorts than the 1998 cohort (3.7% annual growth compared to 2.3% for short higher education graduates and 4.6% compared to 2.4% for long higher education graduates for 1998 and 2010 cohorts). Long higher education graduates suffer most from the slowdown in wage growth: the 2010 cohort's growth loses almost half of its 1998 predecessors' growth. The rest of the analysis will therefore focus on long higher education graduates, although results are presented for all levels of education.

Table 1.6: Log entry wage regressed on number of years spent on the labor market by education level, with gender, location and industry fixed effects

	log entry wage		
	Gen 1998	Gen 2004	Gen 2010
Years × No degree	0.027***	0.037***	0.037***
	(0.003)	(0.003)	(0.003)
Years × High school degree	0.026***	0.019***	0.03***
	(0.001)	(0.001)	(0.002)
Years × Short higher educ. degree	0.037***	0.025***	0.023***
	(0.002)	(0.002)	(0.002)
Years × Long higher educ. degree	0.046***	0.045***	0.024***
	(0.003)	(0.003)	(0.003)
FE education	✓	✓	✓
FE gender	✓	✓	✓
FE location	✓	✓	✓
FE industry	✓	✓	✓
Observations	37 785	27 656	20 130
R^2	0.325	0.244	0.283

*p<0.1; **p<0.05; ***p<0.01

All occupations except Farmers

Estimation results for equation (1.3) are presented in Table 1.7 by cohort, for all levels of education and all occupations, except farmers, because their number is not large enough to obtain a robust estimate. Wage growth heterogeneity between the 1998 and 2010 cohorts by occupation is clearly apparent for higher education graduates. In particular, mid-level managers and top-level managers and highly qualified professionals are particularly affected by the slowdown in wage growth.

Table 1.7: Log entry wage regressed on number of years spent on the labor market by education level and occupation, with gender, location and sector fixed effects

	log entry wage cleaned of fixed effects		
	Gen 1998	Gen 2004	Gen 2010
Years × No degree × Craftmen, Shopkeepers, Business owners	0.075***	0.068***	

	(0.008)	(0.012)	()
Years × No degree × Top managers, Highly qualified prof.	0.064***	0.078***	0.08***
	(0.017)	(0.021)	(0.016)
Years × No degree × Employees	0.016***	0.032***	0.036***
	(0.002)	(0.003)	(0.002)
Years × No degree × Factory workers	0.03***	0.037***	0.037***
	(0.001)	(0.002)	(0.002)
Years × No degree × Mid-level managers	0.022***	0.043***	0.037***
	(0.005)	(0.005)	(0.003)
Years × HS deg. × Craftmen, Shopkeepers, Business owners	0.012***	0.056***	0.022
	(0.003)	(0.004)	(0.024)
Years × HS deg. × Top managers, Highly qualified prof.	0.074***	0.069***	0.068***
	(0.003)	(0.004)	(0.005)
Years × HS deg. × Employees	0.014***	0.006***	0.021***
	(0.001)	(0.001)	(0.001)
Years × HS deg. × Factory workers	0.025***	0.018***	0.026***
	(0.001)	(0.001)	(0.001)
Years × HS deg. × Mid-level managers	0.043***	0.031***	0.041***
	(0.001)	(0.001)	(0.002)
Years × SHE. deg. × Craftmen, Shopkeepers, Business owners	-0.001	0.041***	0.069
	(0.006)	(0.009)	(0.058)
Years × SHE. deg. × Top managers, Highly qualified prof.	0.082***	0.065***	0.065***
	(0.002)	(0.003)	(0.004)
Years × SHE. deg. × Employees	-0.001	-0.013***	-0.004
	(0.002)	(0.002)	(0.003)
Years × SHE. deg. × Factory workers	-0.013***	-0.003	-0.001
	(0.003)	(0.003)	(0.004)
Years × SHE. deg. × Mid-level managers	0.043***	0.034***	0.029***
	(0.001)	(0.001)	(0.002)
Years × LHE deg. × Craftmen, Shopkeepers, Business owners	0.057***	0.048**	0.079**
	(0.009)	(0.019)	(0.035)
Years × LHE deg. × Top managers, Highly qualified prof.	0.058***	0.067***	0.047***
	(0.002)	(0.002)	(0.002)
Years × LHE deg. × Employees	-0.062***	-0.038***	-0.038***
	(0.007)	(0.006)	(0.005)
Years × LHE deg. × Factory workers	-0.045*	-0.049***	-0.054***

	(0.023)	(0.014)	(0.009)
Years × LHE deg. × Mid-level	0.002	0.003	-0.011***
managers	(0.004)	(0.003)	(0.003)
Observations	37 785	27 656	20 130
R^2	0.182	0.139	0.155

*p<0.1; **p<0.05; ***p<0.01

Cells are empty where there were too few observations

All occupations except Farmers

HS: High School, SHE: Short higher educ., LHE: Long higher educ.

Figure 1.2 presents the results of equation (1.5)'s decomposition by level of education. It shows that the slowdown in wage growth for higher education graduates between the 1998 and 2010 cohorts is mainly driven by mid-level managers and top-level managers and highly qualified professionals, who account for almost 100% of the total margin for both short and long higher education graduates. However, the extensive margin of mid-level managers behaves differently from that of top-level managers and highly qualified professionals: it is positive for the former, indicating an increase in the share of long higher education graduates among mid-level managers, and negative for the latter, signalling a decline in their share among top-level managers and highly qualified professionals. This suggests that the increase in the proportion of higher education graduates (especially long higher education) is unmatched by demand of top-level managers and highly qualified professionals. As a result, an increasing share of higher education graduates is absorbed by mid-level managers occupation. Intensive margins for mid-level managers and top-level managers and highly qualified professionals are both negative, for higher education graduates. Because the intensive margin is particularly for short higher education graduates working in mid-level management, it may be that the 'absorption' by mid-level manager of long higher education graduates negatively impacts the growth of entry-level wages. This observation is consistent with a framework of diminishing marginal returns, where last entrants' wages decrease because of their lower productivity.

Figure 1.2: Average wage growth decomposition by education level and occupation

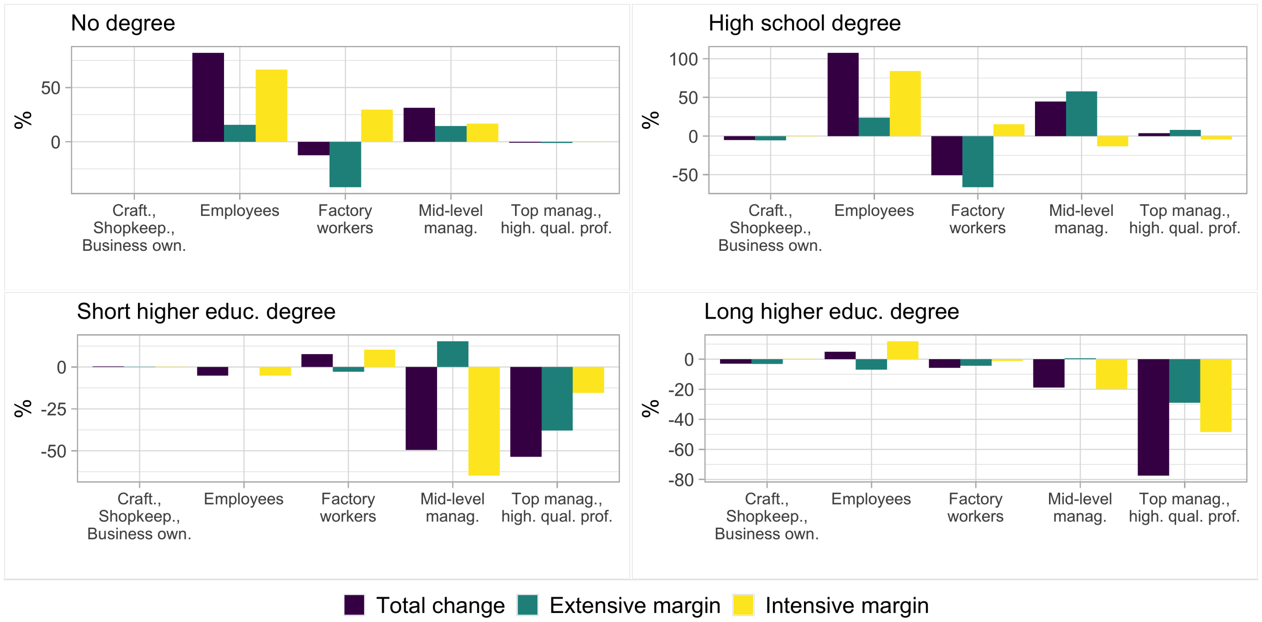

The previous analysis can be replicated at a finer level of occupation aggregation available in the data. The results of this second level of analysis are presented in Figure 1.3, for all higher education graduates and mid-level managers and top-level managers and highly qualified professionals only.

Figure 1.3: Average wage growth decomposition by education level and disaggregated occupation - All higher education graduates

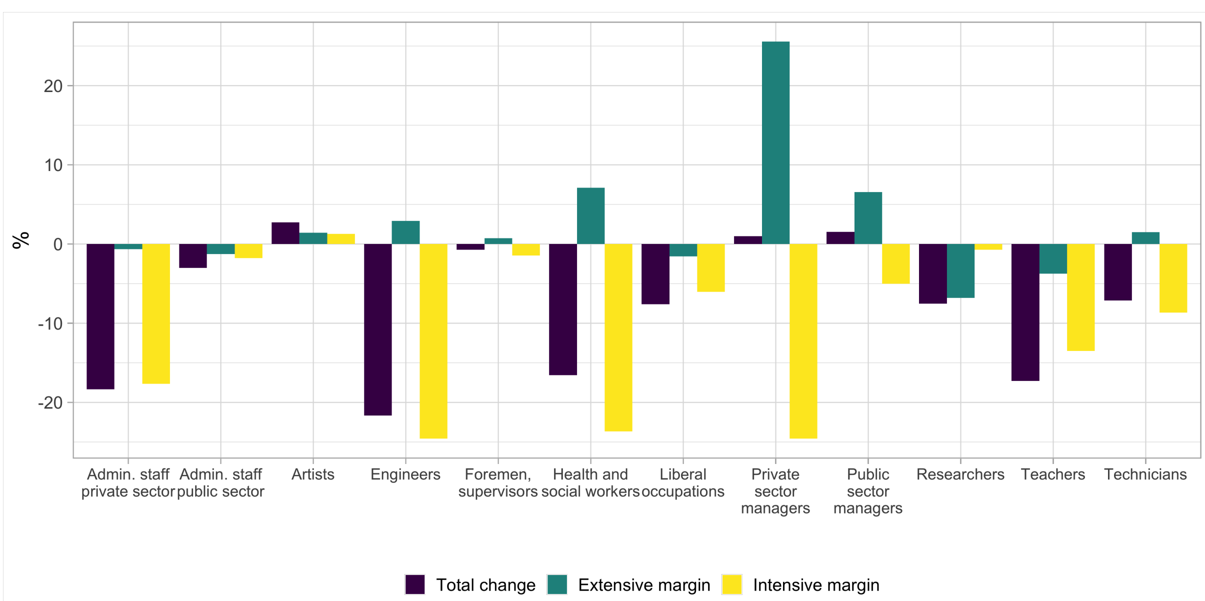

Figure 1.3 shows a significant heterogeneity in margins within the mid-level managers and top-level managers and highly qualified professionals. Sub-occupations can be classified into four categories according to their intensive and extensive margins: the first includes administrative staff in the private sector, teachers, and the liberal occupations. It contributes significantly to the slowdown in wage growth, is driven by both the intensive and extensive margins, although the intensive margin dominates in absolute terms. A second category is made up of engineers and health and social workers, whose contribution to the total slowdown in wage growth at the time of hiring is no less significant but is broken down differently from the first category. Their extensive margin is positive: the proportion of individuals starting contracts in these occupations increased between the 1998 and 2010 cohorts. However, their intensive margin outweighs their extensive margin, and the sum of the two is negative. The third category, which differs from the first two in that it has a positive total margin, comprises public and private sector managers: the contribution of these occupations to wage growth benefits the 2010 cohort compared with the 1998 cohort. This is due solely to a positive extensive margin, which is particularly high for private sector managers. Finally, the last category includes all other sub-occupation that have a small impact on changes in wage growth.

Except for two occupations (researchers and artists), all occupations suffer from a negative intensive margin, i.e. reduced wage growth at hiring, once composition effects (the extensive margin) have been accounted for. Besides, occupations in categories two and three identified above whose intensive margin is largest in absolute value are also those whose extensive margin are most important and positive. This is consistent with a framework of diminishing returns: the influx of new employment contracts in these occupations, driven by an increase in the supply of higher education graduates, leads to a drop in marginal productivity of individuals who have recently entered the market, which translates into lower returns to experience.

In the next section, I explore two possible mechanisms involved in the intensive margin of the wage growth slowdown and how they translate differently into each occupation.

4 Mechanisms

I study two mechanisms that are likely to cause heterogeneity within the wage growth slowdown intensive margin: the first is promotion to managerial positions: it has been documented in the literature (Kwon et al. (2010)) that wage growth is affected if promotions

become scarce for a given cohort. Such a mechanism could interact with the heterogeneity of the wage slowdown by level of education and occupations if the influx of graduates into some occupations, as evidenced by their large extensive margin, is correlated with a decline in the share of managerial positions at hiring, which leads to a lower increase in entry wages. In a theoretical perspective, the reasons for such a correlation are twofold: first, in a framework of decreasing returns an influx of graduates is not followed by demand for managers from firms, which mechanically reduces their proportion among new hires. Second, if the expansion of higher education graduates on the labour market is accompanied by a decline in their unobserved quality, it would lead to a decline in the share of managers within cohorts that experience the educational expansion. To check the validity of this mechanism, I establish two facts: first, lower access to managerial positions is linked to lower wages on hiring, and second, that hiring in those positions does fall for the 2010 cohort compared to the 1998 cohort. If it is particularly the case for occupations with the largest extensive margins , it suggests that the mechanism is rather based on the theory of decreasing returns, since occupations with the largest influx are most affected. If, on the other hand, I observe that the decline managerial positions is the same for all occupations, it would indicate that a decrease in unobserved quality drives the mechanism.

The second mechanism I evidence is the degree specialization (or major) and occupation match quality. In line with theories of human capital, Liu et al. (2016) show on US data that poor match quality in the early years on the labour market weighs on wage developments in the medium to long term. Initial matching can be expected to have an impact on a cohort's medium-term wages in two ways: either initial match quality is the same across cohorts on average, but its impact on subsequent wages becomes stronger and more persistent, or initial match quality decreases across cohorts, and subsequent wage are negatively impacted. I show that the results obtained are consistent with the first explanation.

4.1 Promotion to managerial positions

To determine whether individuals are hired as managers, I use the question 'Do you manage a team?" in the Generations Surveys. This question provides is more accurate for my purpose than using mid-manager or top-manager occupations, as it is unclear whether individuals hired to these occupations do in fact manage some of their colleagues. Table 1.8 shows that the managerial hires increase with the level of education for all cohorts, but also that the share is higher for the long higher education graduates among the 1998 than the 2010 and 2004 cohorts. Besides, for all cohorts and levels of education, the share of managerial hires is higher in the early years on the labour market. Particular attention

should therefore be paid to recruitment opportunities in the first few years on the labour market.

Table 1.8: Share of managerial posisitions obtained by cohort, year 1-2 and year 7-8

	Gen 1998		Gen 2004		Gen 2010	
% Manager	Year 1-2	Year 7-8	Year 1-2	Year 7-8	Year 1-2	Year 7-8
Sans diplôme	19.2	13	11.8	2.2	8.5	12.1
Diplôme du secondaire	17.8	15.6	12.2	12.2	12.8	13.4
Diplôme du tertiaire court	19.8	18	16.4	14.1	18	17.2
Diplôme du tertiaire long	34.6	33.6	28.8	28.1	26	22.9

To understand the link between managerial positions and wage levels in the medium term, I estimate the following regression is made by cohort, only at years 7 and 8:

$$\log w_{ijt} = \sum_{e} \mathbb{1}_{[educ_i=e]} \sum_{p} \mathbb{1}_{[occ_j=p]} \zeta_{gep} \times M_{jt} + g_i + r_j + s_j + \epsilon_{ijt}, \tag{1.6}$$

where M_{jt} is a binary variable equal to 1 if the new job is a manager position, and 0 otherwise. The estimator ζ_{gep} indicates the average wage gain of a manager position compared to a non-manager position in the medium term, by cohort, level of education, and occupation. Long and short higher education graduates are grouped into the same category of higher education graduates.

Table 1.9 presents the three regressions (1.6) for the 1998, 2004 and 2010 cohorts. The impact of a managerial position on entry wages varies according to level of education: significant coefficients are all positive for high school and higher education graduates (except employees for 2004 high school graduates). On the other hand, the relationship is negative and significant for individuals with no degree working as employees and factory workers, indicating that a managerial position at this level of education does not offer the same benefits as to other levels. The relation between managerial position and entry wages is particularly strong for higher education graduates working as mid-level managers, and top managers and highly qualified professionals. However, the intensity of the relationship decreases between the 1998 and 2010 cohorts: among top managers and highly qualified professionals, a managerial position is associated with a salary 63% higher for the 1998 cohort and only 51% higher for the 2010 cohort. This decline could be the result of a drop in managerial productivity due to the particularly large influx of graduates into these occupations. However, the

relationship remains significant and suggests examining the evolution of access to managerial positions between the 1998 and the 2010 cohorts.

Table 1.9: Log entry wage regressed on dummy for managerial position by education level and occupation, with gender, location and sector fixed effects

	log entry wage		
	Gen 1998	Gen 2004	Gen 2010
Manager × No degree× C/S/BO	0.226		
	(0.178)	()	()
Manager × No degree× TM/HQP	0.23		0.231
	(0.303)	()	(0.267)
Manager × No degree× Employees	-0.199**		0.01
	(0.086)	()	(0.084)
Manager × No degree× Factory workers	-0.103*	0.094	-0.143***
	(0.055)	(0.177)	(0.053)
Manager × No degree× Mid-level managers	0.067	0.574**	0.111
	(0.149)	(0.228)	(0.083)
Manager × HS deg.× C/S/BO	-0.043		
	(0.078)	()	()
Manager × HS deg.× TM/HQP	0.554***	0.374***	0.374***
	(0.086)	(0.086)	(0.078)
Manager × HS deg.× Employees	-0.048	-0.105**	-0.052
	(0.04)	(0.051)	(0.065)
Manager × HS deg.× Factory workers	0.045	0.073*	0.023
	(0.032)	(0.038)	(0.054)
Manager × HS deg.× Mid-level managers	0.199***	0.155***	0.181***
	(0.035)	(0.034)	(0.041)
Manager × HE. deg.× C/S/BO	0.575***		0.568
	(0.12)	()	(0.378)
Manager × HE. deg.× TM/HQP	0.634***	0.622***	0.508***
	(0.028)	(0.038)	(0.037)
Manager × HE. deg.× Employees	0.101	0.21**	0.148
	(0.107)	(0.103)	(0.117)
Manager × HE. deg.× Factory workers	0.013	0.222*	0.073
	(0.118)	(0.114)	(0.114)
Manager × HE. deg.× Mid-level managers	0.413***	0.248***	0.354***

	(0.043)	(0.043)	(0.048)
FE gender, location, industry	✓	✓	✓
Observations	4 730	3 433	2 792
R^2	0.332	0.277	0.291

*p<0.1; **p<0.05; ***p<0.01
Cells are empty where there were too few observations
Observations are the sequences between years 7 and 8
HS: High School, HE: Higher educ.
C/S/BO: Craftmen, Shopkeepers, Business owners
TM/HQP: Top managers, Highly qualified prof.

Probability to access a manager position by cohort and level of education within each occupation is assessed using the following logistic regression:

$$M_{jt} = \sum_{e} \mathbb{1}_{[educ_i=e]} \sum_{p} \mathbb{1}_{[occ_j=p]} \pi_{gep} \times a_t + e_i + o_j + g_i + r_j + s_j + \epsilon_{ijt}, \tag{1.7}$$

where M_{jt} is the binary variable equal to 1 if the individual is hired as a manager and 0 otherwise. Years since leaving the education system a_t range from 1 to 6 inclusive. The coefficient $\hat{p}_{gep} = 100 \times (\exp(\hat{\pi}_{gep}) - 1)$ informs on the probability increase each year in percentage terms.

Table 1.10 shows the change in $\hat{p}_{gep}$ between 1998 and 2010, by level of education. Among 1998 higher education graduates hired a top managers or highly qualified professional, chances of accessing a managerial position increase by about 1.2% per year. This percentage is negative and non-significant for the 2010 cohort. For middle-managers the increase in probability is non significant for the 1998 cohort, and negative and significant for the 2010 cohort (-1.1%).

Table 1.10: Chances of obtaining a managerial postition by years spent on the labor market, by education level and occupation

	Managing position		
	Gen 1998	Gen 2004	Gen 2010
Year × No degree × C/S/BO	3.968* (0.02)	()	()

Year × No degree × TM/HQP	17.697***	-3.971	-5.156
	(0.026)	(0.058)	(0.042)
Year × No degree × Employees	-1.464***	-1.212*	0.921*
	(0.006)	(0.007)	(0.005)
Year × No degree × Factory workers	-0.876*	-1.14*	0.519
	(0.005)	(0.006)	(0.005)
Year × No degree × Mid-level managers	0.751	-4.375***	1.142
	(0.009)	(0.009)	(0.008)
Year × HS deg. × C/S/BO	0.06	3.692	0.521
	(0.014)	(0.074)	(0.047)
Year × HS deg. × TM/HQP	3.281***	-1.185	-1.555*
	(0.007)	(0.008)	(0.009)
Year × HS deg. × Employees	-0.641**	-0.012	-1.054***
	(0.003)	(0.003)	(0.003)
Year × HS deg. × Factory workers	-0.305	0.309	-0.207
	(0.003)	(0.003)	(0.004)
Year × HS deg. × Mid-level managers	-0.126	-0.515	-0.571
	(0.004)	(0.004)	(0.004)
Year × HE. deg. × C/S/BO	2.43*	13.451**	2.604
	(0.014)	(0.05)	(0.065)
Year × HE. deg. × TM/HQP	1.208***	1.237***	-0.492
	(0.004)	(0.004)	(0.004)
Year × HE. deg. × Employees	-1.934***	-0.613	-0.995*
	(0.004)	(0.004)	(0.005)
Year × HE. deg. × Factory workers	-0.524	0.827	0.43
	(0.007)	(0.006)	(0.007)
Year × HE. deg. × Mid-level managers	0.123	0.642*	-1.137***
	(0.003)	(0.003)	(0.004)
FE education, occupation	✓	✓	✓
FE gender, location, industry	✓	✓	✓
Observations	32 700	23 226	17 301
R^2	0.059	0.071	0.066

*p<0.1; **p<0.05; ***p<0.01

Cells are empty where there were too few observations

Observations are the sequences between years 1 and 6

HS: High School, HE: higher educ.

TM/HQP: Top managers, Highly qualified prof.

C/S/BO: Craftmen, Shopkeepers, Business owners

4.2 Major-Occupation match quality

Match quality between degree and employment plays an important role in the persistence of initial economic conditions' effect on medium term wage (Liu et al. (2016)). By defining match quality by wage level in first job for each degree specialization in each industry, the authors establish that the deterioration of initial match quality match due to poor economic conditions in the United States in the 2010s (including high unemployment) has a downward impact on wage levels in subsequent jobs. This mechanism is particularly salient for higher education graduates, who are more specialized than their less educated peers. In France, in addition to the recession of the early 2010s, there has been an increase and diversification of the educational supply (Dupray and Moullet (2010)) among the 2004 and 2010 cohorts, which may also have affected the quality of initial matching the first job for these two cohorts.

Table 1.11 shows the list of degree majors available to higher education students in France and their distribution among each cohort. Despite the diversification of educational provision, their distribution remains fairly stable between the 1998 and 2010 cohorts. There is however a decline in the 'Mechanics, Electricity' and 'Electronics' specialization and the rise of the 'Mathematics and Science' and 'Social work' specializations.

Table 1.11: Degree specialization shares by cohort

Share of graduates (%)	Gen 1998	Gen 2004	Gen 2010
Agriculture, fishing & woodland	4.4	4.4	3.8
Civil engineering & Construction	4.7	5.1	4.3
Communication & information	7	6.8	4.8
Community services	2.5	1.5	2
Flexible materials	1	0.9	0.5
General production	3.6	4.8	6
General service	0.2	0.5	3.3
General training	4.4		0.5
Humanities & law	10.8	8.3	9
Industrial transformations	5.6	5.9	4.9
Literature & arts	4.7	4.7	5.1
Mathematics & sciences	3.7	4.9	9.1
Mechanics, electricty & electronics	13.7	12.1	7.9
Personal services	15.7	18.9	19.7
Trade & management	17.9	21.3	19
Total	100	100	100

Thanks to the detailed level of data, I can analyse the specialization-occupation match quality. It is defined by the following regression, performed only on contracts starting in the year in which each cohort enters the labour market, by cohort and education level:

$$\log w_{ijt} = \sum_s \mathbb{1}_{[spec_i=s]} \sum_p \mathbb{1}_{[occ_j=p]} \delta_{geps} + \epsilon_{ijt}, \tag{1.8}$$

where $spec_i is$ the specialty chosen by individual i during their studies. The estimated coefficient δ_{geps} is an average of the logarithm of the first year's wage on the labour market, by cohort, education level, occupation, and degree specilization. To define a measure of matching, I look for the best matched specialization within a cohort, education level and occupation, i.e. the one for which average wage is highest:

$$\delta^*_{gep} = \max_s \delta_{geps}. \tag{1.9}$$

Matching quality for a given degree specialization is defined by how far it stands with respect to the best matched specialization:

$$D_{geps} = -\left|\delta_{geps} - \delta^*_{geps}\right|. \tag{1.10}$$

D_{geps} is always below or equal to zero, if s is the best matched specialization. The difference between δ_{geps} and $\delta_{geps'}$ for two specialisations s and s' is interpreted as the percentage difference between the average hiring wage for s and for s'. The higher the absolute value of D_{geps}, the farther average salary for the specialization s is from best matched specialization. Matching quality is then be said to be poor. D_{geps} is a flexible measure of matching since one specialization may be mismatched with one occupation, but well matched for another. Matching quality is computed at cohort and education level, so that average earnings comparison between cohorts are irrelevant to computing D_{geps}.

Individuals who are not hired in their first year on the labour market are not included in regression (1.8). These individuals are excluded from the analysis, which therefore covers only a subset of each cohort. Another definition for the period during which initial matching quality is computed is explored in the robustness tests.

Table 1.12 shows the evolution of measure D_{geps} by cohort and education level in terms of median and interquartile deviation, weighted by individuals. Since individuals with no degree do not choose a specialization, they are excluded from the present analysis. Although median quality of matching deteriorates for high school graduates, it increases for higher education graduates, while the interquartile gap is the same for higher education graduates between 1998 and 2010 cohorts (and narrower for the 2004 cohort). Worsening of matching quality does not therefore appear to be a factor in wage growth slowdown for higher education graduates. It remains to be determined whether, despite the consistent quality of matching across the three cohorts, its effect on the persistence of wage levels changed between the 1998 and 2010 cohort.

Table 1.12: Match quality: median and interquartile range by cohort and education level

	Education level	Gen 1998	Gen 2004	Gen 2010
High school degree	p50	-0.2	-0.16	-0.22
	[p25-p75]	[-0.33,-0.09]	[-0.26,-0.09]	[-0.38,-0.12]
Short higher educ. degree	p50	-0.23	-0.16	-0.21
	[p25-p75]	[-0.32,-0.1]	[-0.31,-0.05]	[-0.31,-0.09]
Long higher educ. degree	p50	-0.38	-0.13	-0.23
	[p25-p75]	[-0.46,-0.16]	[-0.26,-0.07]	[-0.32,-0.13]

The impact of initial match quality on entry wages in subsequent years is assessed by the following regression, at cohort and education level:

$$\log w_{ijt} = \sum_a \mathbb{1}_{[year_t=a]} D_{geps} \lambda_{gea} + a_t + g_i + r_j + s_j + \epsilon_{ijt}. \tag{1.11}$$

Unlike the previous regressions, the aim is to obtain an estimated coefficient $\hat{\lambda}_{geps}$ differentiated per year. Within a cohort and education level, all individuals face the same conditions every year, captured by a fixed effect, hence the only variation between individuals in this regression is due to the difference in initial matching quality.

The estimated coefficients $\hat{\lambda}_{geps}$ are presented in Table 1.13 for long higher education graduates only. Distinguishing between short and long higher education graduates matters in this analysis because degree specialization is closely linked to education level. Table 1.13 shows significant intergenerational differences among long higher education graduates: in the first years on the labour market, initial match quality's effect on entry wages is similar for all cohort. In year 1, a 1 percentage point increase in match quality, i.e. a .01 relative increase of $\hat{\delta}_{geps}$ over $\hat{\delta}^*_{geps}$, results in a wage increase of almost 1% for all generations (.72%, .83% and .67% respectively). However the effect of initial match quality diverge between cohorts around year 4, since they are no longer significant for the 1998 cohort, whereas they persist until year 7 for the 2004 cohort and until year 8 for the 2010 cohort.

Table 1.13: Log entry wage regressed on match quality by year and education level - Long higher education graduates

	log entry wage		
	Gen 1998	Gen 2004	Gen 2010
Year 1 × match quality	0.718***	0.828***	0.665***
	(0.057)	(0.052)	(0.034)
Year 2 × match quality	0.264***	0.223**	0.342***
	(0.099)	(0.092)	(0.058)
Year 3 × match quality	0.424***	0.481***	-0.012
	(0.126)	(0.096)	(0.066)
Year 4 × match quality	0.012	0.198**	0.272***
	(0.136)	(0.093)	(0.074)
Year 5 × match quality	-0.177	0.575***	0.363***
	(0.108)	(0.122)	(0.084)
Year 6 × match quality	0.087	0.343**	0.087
	(0.128)	(0.166)	(0.093)
Year 7 × match quality	-0.029	0.415***	0.297***
	(0.159)	(0.119)	(0.086)
Year 8 × match quality	0.285*	-0.123	0.195**
	(0.166)	(0.161)	(0.079)
FE gender, location, industry	✓	✓	✓
Observations	1 634	2 550	4 298
R^2	0.43	0.325	0.291

*p<0.1; **p<0.05; ***p<0.01

Only individuals who found a job in the first year after leaving school are included

Initial match quality plays a significant role in slowing down the salary progression of long higher education graduates between the 1998 and 2010 cohorts, not because the quality has deteriorated, but because its impact on subsequent salary levels has increased. This could be explained by mobility: if the 2010 cohort would change jobs less frequently than the 1998 cohort, the initial match quality may play a role in determining hiring wages for a longer time. However, average number of employment spells for long higher education graduates is 4.1,

4.5 and 4.4 for cohorts 1998, 2004 and 2010 respectively. The two last cohorts are changing jobs as much, if not more, than the 1998 cohort. To explore the mobility hypothesis in more detail, it is necessary to analyse the structure of transitions between occupations of each generation. If the 2004 and 2010 cohorts switch occupations less often than the 1998 cohort, initial matching quality's impact may increase as individuals remain stuck in occupations to which their degree is not adapted. However, this hypothesis is again contradicted by the data: among the long tertiary graduates, those who never change occupation in the first seven years on the labour market account for 40%, 34% and 32% of the 1998, 2004 and 2010 cohorts, respectively. Individuals changing only once accounted for 31%, 29% and 30% of each cohort. This shows changes in occupations are more common among the 2004 and 2010 cohorts. Finally, the increased impact of initial match quality could be explained by the rise in the number of graduates in the labour market: mismatched individuals in the 2004 and 2010 cohorts may find it harder to access higher-paying jobs for their degree specialization as each year a new cohort enters the labour market, increasing competition for the best jobs. Because the supply of long higher education graduates is smaller in the late 1990s and early 2000s, the 1998 generation faces less competition and is able to make up for any low initial match quality.

5 Robustness tests

5.1 Sample representativity

The Generation surveys only provide information on wages when individuals transition from a job to another, or when they transition from and to unemployment. As such, they constitute an unbalanced panel: some individuals are not observed in some years. An individual that transitions often makes up more observations than an individual who stays in the same spell over the period, and thus have a greater weight in the data. It is therefore important to check that individuals that go through few transitions, such as those hired in their first year on the labour market who remain in their jobs for the next seven years, experience the same trend in wage growth between the generations 1998 and 2010. To do this, I perform two analyses: the first uses the wages observed during the last interview session of each generation. If the interviewee is employed during the last session, his or her current salary is reported as exit wage, even if his or her employment spell is not ending. This provides a cross-section of the entire population surveyed at the end of 2005, 2011 and 2017 respectively for each of the 1998, 2004 and 2010 cohorts. Comparing wages in this cross-section by education level and across cohort is a way of checking all graduates are affected

by declining returns to experience. The second analysis looks at annual average exit wages over the seven years and examines whether entry wage growth slowdown is compensated for by pay raise during employment spells.

Table 1.14 shows the average observed wages in the 2005, 2011 and 2017 cross-sections (in constant euro base 2017) by education level and cohort. These cross-sectional wages have decreased on average between the 1998 and 2010 cohort for higher education graduates, which confirms that wage growth has slowed for all individuals in the cohort, including those in long-term employment.

Table 1.14: Average observed wage at end of survey, by cohort and education level

Niveau Education	Gen 1998	Gen 2004	Gen 2010
Sans diplôme	1325	1341	1357
Diplôme du secondaire	1499	1472	1508
Diplôme du tertiaire court	1918	1775	1826
Diplôme du tertiaire long	2902	2594	2567

Finally, Figure 1.4 shows changes in average exit wages (i.e. the last wage received in the job) over time by education level and cohorts. These wages are higher than hiring wages for all cohorts and education levels but exhibit the same slowdown trend for the higher education graduates between 1998 and 2010 as entry wages. I conclude that pay raises on the job do not compensate for decreasing returns to experience at hiring.

Figure 1.4: Average real wage at employment spell's end, by cohort and education level

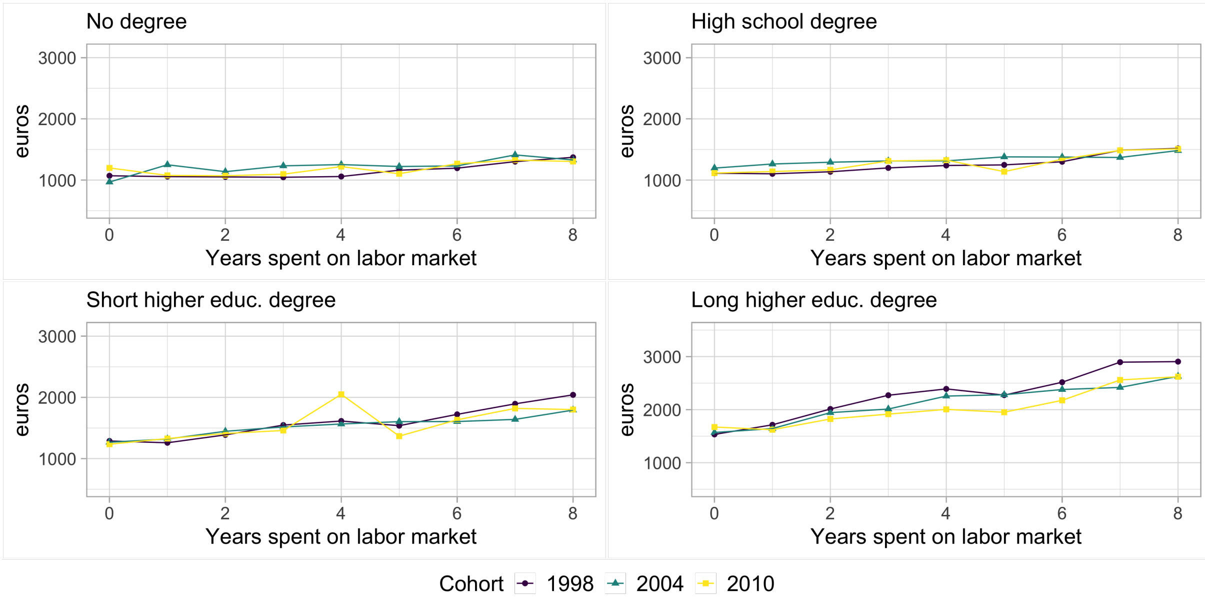

5.2 Unobserved heterogeneity

The results presented above are based on the identification assumption that the distribution of unobserved quality is there within each cohort. I test this assumption by using a proxy for unobserved quality, which is grade repetition before the start of secondary school. The Generation surveys provide individuals' age in 6th grade (the first grade in secondary school). Normal age is 6th grade is 11 years old, hence if an individual is older when entering 6th grade, I deduce they have repeated a grade in primary school. Repeating a grade before 6th grade indicates lower academic and learning abilities, which in turn affects the individual's wage levels in the labour market. I am agnostic as to the causes of these lower abilities. 23.0% of individuals in the 1998 cohort, 12.1% of individuals in the 2004 cohort and 12.8% of individuals in the 2010 cohort repeated a grade in primary school. The practise of grade repeating scaled back over the period, hence the high number of individuals who repeat grade in the 1998 generation is only partly indicative of a lower average unobserved quality.

To check the impact of unobserved quality on wage levels I introduce a dummy for class repetition in my baseline regression:

$$\log w_{ijt} = \sum_{e} \mathbb{1}_{[educ_i=e]} \times a_t + \alpha_i + g_i + r_j + s_j + \epsilon_{ijt}, \tag{1.12}$$

where α_i is equal to 1 if the individual has repeated a grade before entering secondary school and 0 otherwise. Since class repetition is only an imperfect measure of unobserved quality, and its practise has evolved between the 1998 and 2010 cohort, I do not compare its effect on log wages between cohorts. Instead, $\hat{\beta}_{eg}$ will be useful to understand if the effect observed in the baseline analysis is solely due to variation in unobserved quality.

Regression estimation is presented in Table 1.15. Grade repetition has a significant and negative effect on wage for all cohorts. It does not significantly change the previous results however: the slowdown in wage growth for long and short higher education graduates between 1998 and 2010 remains qualitatively the same. This suggests that this slowdown is not due to unobserved quality variations between the 1998 and 2010 cohorts.

Table 1.15: Log entry wage regressed on dummy for grade repeat and years spent on the labor market by education level

	log entry wage		
	Gen 1998	Gen 2004	Gen 2010
Grade repeat	-0.06***	-0.046***	-0.058***
	(0.004)	(0.006)	(0.007)
Years × Sans diplôme	-0.01***	-0.009***	-0.007***
	(0.001)	(0.002)	(0.002)
Years × High school degree	0.01***	0.007***	0.01***
	(0.001)	(0.001)	(0.001)
Years × Short higher educ. degree	0.054***	0.038***	0.04***
	(0.001)	(0.001)	(0.002)
Years × Long higher educ. degree	0.117***	0.099***	0.084***
	(0.002)	(0.002)	(0.002)
FE gender, location, industry	✓	✓	✓
Observations	37 785	27 599	19 992
R^2	0.29	0.218	0.229

*p<0.1; **p<0.05; ***p<0.01

Only individuals whose age is known in 6th grade are included

A second measure of unobserved quality, specific to long higher education graduates, is

the type of school in which individuals obtained their degree, and whether they obtained a master's degree or PhD. There are distinction in France between public university, where students can graduate both from a master and a PhD, and "Grandes Ecoles", specific schools specialized in engineering or business, that deliver specific degrees. Engineering and business schools graduates often obtain higher wages than university graduates. Since business and engineering schools have traditionally been more selective than universities, it can be expected that they will continue to be selective even if more young people apply. In universities, on the other hand, there is no selection at entry since any student can enrol for a bachelor's degree (except in a few courses where seats are scarce). It may then be that the 1998 cohort count proportionally more engineering and business school graduates, driving average wages upwards. Table 1.16 shows the distribution of types of degrees obtained among long higher education graduates by cohorts. PhD are not accounted for among the 1998 cohort and the type of degree is not unknown for 1.5% of graduates in the 2004 cohort. The share of university graduates (master's and doctoral degrees) indeed increased between cohorts 1998 and 2010.

Table 1.16: Degree type shares among long higher education graduates

Degree type (%)	Gen 1998	Gen 2004	Gen 2010
Business degree	10.1	11.1	9.2
Engineering degree	25.5	23.4	19.9
Masters degree	64.4	47.3	58.8
Doctorat		16.6	12.1
Inconnu		1.5	

To check that the wage growth slowdown is not due to a composition effect on the type of schools between the 1998 and 2010 cohorts, I perform the following regression, only for graduates of the long higher education sector:

$$\log w_{ijt}^0 = \sum_d \mathbb{1}_{[degree_i = d]} \xi_{gd} \times a_t + + \epsilon_{ijt}, \tag{1.13}$$

where $\log w_{it}^0$ is log wage cleaned of fixed effects for gender, region and industry that was computed in section 2 equation (1.2). Estimate $\hat{\xi}_{gd}$ capture wage growth by degree type. It is presented in Table 1.17. For Individuals who graduated from business and engineering schools, where the selection should have remained stronger than at university, wage growth slowed as much as at university (-42, -33, -33 percentage points for business, engineering, and master's graduates, respectively, between generations 1998 and 2010). The type of

degree does not capture unobserved quality that could be driving the wage growth slowdown observed between cohorts.

Table 1.17: Log entry wage regressed on years spent on the labor market by type of degree

	log entry wage		
	Gen 1998	Gen 2004	Gen 2010
Years × Business degree	0.081***	0.073***	0.038***
	(0.005)	(0.004)	(0.004)
Years × Engineering degree	0.065***	0.057***	0.042***
	(0.003)	(0.004)	(0.003)
Years × Masters degree	0.032***	0.026***	0.011***
	(0.002)	(0.002)	(0.002)
Years × Doctorat		0.071***	0.063***
	()	(0.004)	(0.004)
FE gender, location, industry	✓	✓	✓
Observations	2 787	3 835	5 891
R^2	0.105	0.091	0.053

*p<0.1; **p<0.05; ***p<0.01

Individuals whose degree type is unknown are excluded

PhDs are not accounted for in Gen 1998

5.3 Matching quality

Finally, I test for a different period over which initial matching quality is computed. In the baseline, the period considered is the first year on the labour market, i.e. the first observation year of the entire cohort (i.e. 1998, 2004 and 2010 respectively). However, this definition is restrictive in that it does not always leave a full year to everyone. For example, if an individual graduates in June, it leaves only 6 months, between June and December, to observe a first hire. The benefit of the baseline definition is to ensure that individuals face the same conditions in the labour market over a limited period, but it may neglect first hires for individuals finding their first job early in the year after the entire cohort leaves the education system. The alternative period definition considers the first year on the labour market at individual rather than cohort level: for instance, if an individual graduates in

June 2010, their first year on the labour market runs from July 2010 to June 2011. I then carry out the same analysis to study the impact of initial matching quality on wage growth.

The results of the analysis using modified matching quality are presented in Table **??**. These results are qualitatively similar to the baseline: the 2010 cohort experiences a longer impact of initial matching quality on wage levels than the 1998 generation. However, the effect is shorter for cohorts 2004 and 2010, since it is no longer significant from years 4 and 6 (compared to 7 and 8 in the baseline analysis). R^2 are higher in the reference regression for the 1998 and 2004 cohorts, and almost equal for the 2010 cohort despite the increase in the number of observations between the reference regression and this one. Hence the shorter effect may be due to a loss in precision due to the new definition of initial period, that doesn't hold initial job market conditions constant.

Table 1.18: Log entry wage regressed on alternative match quality by year and education level - Long higher education graduates

	log entry wage		
	Gen 1998	Gen 2004	Gen 2010
Year 1 × match quality	0.944***	0.864***	0.62***
	(0.048)	(0.044)	(0.025)
Year 2 × match quality	0.529***	0.486***	0.315***
	(0.061)	(0.058)	(0.033)
Year 3 × match quality	0.488***	0.501***	0.222***
	(0.071)	(0.074)	(0.041)
Year 4 × match quality	0.09	-0.003	0.236***
	(0.092)	(0.07)	(0.046)
Year 5 × match quality	-0.1	0.041	0.241***
	(0.093)	(0.086)	(0.056)
Year 6 × match quality	-0.029	0.042	0.089
	(0.099)	(0.11)	(0.059)
Year 7 × match quality	0.101	-0.249***	-0.022
	(0.102)	(0.096)	(0.055)
Year 8 × match quality	-0.019	-0.538***	0.018
	(0.109)	(0.111)	(0.054)
FE gender, location, industry	✓	✓	✓
Observations	5 559	4 721	5 116
R^2	0.188	0.202	0.3

*p<0.1; **p<0.05; ***p<0.01

Only individuals who found a job in the first year after leaving school are included

6 Conclusion

The Generations Surveys highlight a delay in wage growth that affect cohorts leaving the higher education system in 2004 and 2010 compared to the cohort who graduates in 1998. I decompose the wage growth slowdown by occupation in two margins: an extensive margin, which reflects changes in the distribution of occupations within each cohort, and an

intensive margin, that captures changes in hiring wages by occupation between cohorts. I perform the decomposition at a finer level of occupation and find a clear heterogeneity in the extensive margin between occupations in middle management, and top management and highly qualified professionals. Among these two categories, occupations which experience the largest influx of higher education graduates between the 1998 and 2010 cohorts are also those for which the intensive margin is largest. It suggests the influx of young graduates has not increased the productivity of companies as much as their senior counterparts, which has impacted their wage growth downward. I further explore possible mechanisms in line with this interpretation: access to manager positions and initial match quality. I find both mechanisms play a role in the young higher education graduates' wage growth slowdown.

References

Ashworth, J., Hotz, V. J., Maurel, A., and Ransom, T. (2021). Changes across Cohorts in Wage Returns to Schooling and Early Work Experiences. *Journal of Labor Economics*, 39(4):931–964.

Beaudry, P., Green, D. A., and Sand, B. M. (2015). The Great Reversal in the Demand for Skill and Cognitive Tasks. *Journal of Labor Economics*, 34(S1):S199–S247.

Berge, W. v. d. (2018). Bad start, bad match? The early career effects of graduating in a recession for vocational and academic graduates. *Labour Economics*, 53:75–96.

Berge, W. v. d. and Brouwers, A. (2017). A lost generation? The early career effects of graduating during a recession. Technical report, CPB Netherlands Bureau for Economic Policy Analysis.

Blair, P. Q. and Deming, D. J. (2020). Structural Increases in Demand for Skill after the Great Recession. *AEA Papers and Proceedings*, 110:362–365.

Blundell, R., Green, D. A., and Jin, W. (2022). The U.K. as a Technological Follower: Higher Education Expansion and the College Wage Premium. *The Review of Economic Studies*, 89(1):142–180.

Bound, J. and Johnson, G. (1992). Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations. *The American Economic Review*, 82(3):371–392.

Brunner, B. and Kuhn, A. (2014). The impact of labor market entry conditions on initial job assignment and wages. *Journal of Population Economics*, 27(3):705–738.

Card, D. and DiNardo, J. E. (2002). Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles. *Journal of Labor Economics*, 20(4):733–783.

Card, D. and Lemieux, T. (2001). Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis. *The Quarterly Journal of Economics*, 116(2):705–746.

Carneiro, P. and Lee, S. (2011). Trends in Quality-Adjusted Skill Premia in the United States, 1960-2000. *American Economic Review*, 101(6):2309–2349.

Doepke, M. and Gaetani, R. (2020). Why Didn't the College Premium Rise Everywhere? Employment Protection and On-the-Job Investment in Skills. *IZA Working Paper*.

Dupray, A. and Moullet, S. (2010). Offre éducative, valorisation des diplômes et effets de composition : deux générations de sortants de l'université au tournant des années 2000. *Revue d'economie politique*, Vol. 120(5):845–880.

Epiphane, D., Mazari, Z., Olaria, M., and Sulzer, E. (2019). Des débuts de carrière plus chaotiques pour une génération plus diplômée. CEREQ Bref 382.

Fredriksson, P., Hensvik, L., and Skans, O. N. (2018). Mismatch of Talent: Evidence on Match Quality, Entry Wages, and Job Mobility. *American Economic Review*, 108(11):3303–3338.

Gaini, M., Leduc, A., and Vicard, A. (2013). Peut on parler de générations sacrifiées ? Entrer sur le marché du travail dans une période de mauvaise conjoncture économique. *Économie et Statistique*, n 462-463.

Gregg, P. and Tominey, E. (2005). The wage scar from male youth unemployment. *Labour Economics*, 12(4):487–509.

Guvenen, F., Kaplan, G., Song, J., and Weidner, J. (2017). Lifetime Incomes in the United States over Six Decades. NBER Working Paper 23371.

Hershbein, B. and Kahn, L. B. (2018). Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings. *American Economic Review*, 108(7):1737–1772.

Jeong, H., Kim, Y., and Manovskii, I. (2015). The Price of Experience. *American Economic Review*, 105(2):784–815.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Kong, Y. C., Ravikumar, B., and Vandenbroucke, G. (2018). Explaining cross-cohort differences in life-cycle earnings. *European Economic Review*, 107:157–184.

Kwon, I., Milgrom, E. M., and Hwang, S. (2010). Cohort Effects in Promotions and Wages: Evidence from Sweden and the United States. *The Journal of Human Resources*, 45(3):772–808.

Lindley, J. and Machin, S. (2016). The Rising Postgraduate Wage Premium. *Economica*, 83(330):281–306.

Lindley, J. and McIntosh, S. (2015). Growth in within graduate wage inequality: The role of subjects, cognitive skill dispersion and occupational concentration. *Labour Economics*, 37:101–111.

Liu, K., Salvanes, K. G., and Sørensen, E. O. (2016). Good skills in bad times: Cyclical skill mismatch and the long-term effects of graduating in a recession. *European Economic Review*, 84:3–17.

Manovskii, I. and Kambourov, G. (2005). Accounting for the Changing Life-Cycle Profile of Earnings. 2005 Meeting Paper 231, Society for Economic Dynamics.

Oreopoulos, P., von Wachter, T., and Heisz, A. (2012). The Short- and Long-Term Career Effects of Graduating in a Recession. *American Economic Journal: Applied Economics*, 4(1):1–29.

Rothstein, J. (2020). The Lost Generation? Labor Market Outcomes for Post Great Recession Entrants. NBER Working Paper 27516.

Stevens, K. (2008). Adverse Economic Conditions at Labour Market Entry: Permanent Scars or Rapid Catch-up. *Working Paper*.

Valletta, R. G. (2016). Recent Flattening in the Higher Education Wage Premium: Polarization, Skill Downgrading, or Both? NBER Working Paper 22935.

Verdugo, G. (2014). The great compression of the French wage structure, 1969–2008. *Labour Economics*, 28:131–144.

von Wachter, T. (2020). The Persistent Effects of Initial Labor Market Conditions for Young Adults and Their Sources. *Journal of Economic Perspectives*, 34(4):168–194.

CHAPTER 2

Education Expansion, Sorting, and the Decreasing Education Wage Premium

Abstract

This chapter studies the interplay between worker supply and firm demand, and their effect on sorting and wages in the labor market. I build a model of one-to-many matching with multidimensional types in which several workers are employed by a single firm. Matching is dictated by worker preferences, their relative productivity in the firm, and substitution patterns with other workers. Using tools from the optimal transport literature, I solve the model and structurally estimate it on Portuguese matched employer-employee data. The Portuguese labor market is characterized by an increase in the relative supply of high school graduates, an increasingly unbalanced distribution of high school graduates versus non-graduates across industries, and a decreasing high school wage premium between 1987 and 2017. Counterfactual exercises suggest that both changes in worker preferences and the increasing relative productivity of high school graduates over non-graduates act as a mitigating force on the decreasing high school wage premium, but do not fully compensate for high school graduates' rise in relative supply.

1 Introduction

Between the 1970s and today, many economies both in the developed and developing world have experienced an increase in their educated labor supply. As a result, the ratio of educated workers (whether high school or college-educated) to uneducated workers present in labor markets has risen. The shift in labor supply's education level has induced broad changes in labor markets, both in terms of workers' allocation to firms and wage structure. Workers-firm allocation, or matching, refers in this chapter to how workers of different education and experience level sort with firms in different sectors. It generates dispersion in the wage structure through its impact on the output produced by a firm and its workforce. The output is determined both by the worker-firm match and the match between a worker and his co-workers. In particular, workers-firm allocation has repercussions on the wage returns to education. This paper seeks to provide a theoretical framework to understand how changes in labor supply affects matching between workers and firms, and through this channel, impacts returns to education. It proposes a novel model of matching on the labor market in which a single firm matches with several workers. The model is structurally estimated on Portuguese matched employer-employee data. In doing so, I am able to quantify the impact of supply and demand changes on worker-firm allocation and wage structure.

The mechanisms driving matching between workers and firms and the resulting wage distribution are two-sided. On the one hand employed workers with various education and experience interact within the firm to produce an output, whose level depends on a production function that is particular to each firm's sector. Given their production function, firms seek to hire a workforce, which is a mix of workers with different characteristics, to maximize their profit. On the other hand, workers have preferences for the tasks performed on the job, which vary from one sector to another. Worker preferences impact which sector they are willing to work in. Given distributions of education and experience in the worker population and sectoral composition among firms, firm production requirements and worker preferences result in a given level of sorting and wage gaps. Sorting is the result of worker-firm allocation: it is the ratio of educated to non-educated workers in each sector. Wage gaps summarize the wage distribution: they are the ratio of educated workers' to non-educated workers' average wage.

To capture these mechanisms, I build a static one-to-many matching model with transferable utility. Workers and firms differ with respect to their observed characteristics, which are summarized by a multidimensional type, as well as a stochastic shock that accounts

for unobserved heterogeneity. A single firm matches with several workers, who constitute a bundle that forms its workforce. The surplus created by the match depends on the firms' observable characteristics as well as the workforce. The utility is transferable under the form of wages paid by the firm to the workers in its workforce. Firms seek to maximize total profit, which is additive in the difference of production and total wage bill, plus random shocks. Workers maximize their utility, which is additive in amenities, wage, and a random shock. Amenities embody workers' inner preference for a given type of firm. At equilibrium, wages clear the market and each agent matches with their best option given wages. The model can generate a rich distribution of wages that depend both on workers' and firms' observable characteristics, as well as on the employed workforce. It also predicts equilibrium matching, which is the joint distribution of firms and workforces. Using both matching and wages, I can separately identify firm production from workers' amenities.

The framework offers more flexibility in estimation than classic supply and demand models developed in Katz and Murphy (1992) and Card and Lemieux (2001): it identifies worker preferences in addition to firm production, as well as varying production parameters over time, which allows for non-linearities in the evolution of firm production paramaters. This is because by explicitly modeling firms' and workers' match choices, I can use both observed matching and observed wages, which brings more power to identification. The model is fitted to the data by assuming parametric forms for firm production and workers' amenities. I classify workers into two education levels, high school graduates and non-graduates, and three age groups, young, middle-aged, and senior. Firms are differentiated by their sector of activity. Following the literature, I choose a nested Constant Elasticity of Substitution (CES) function for production, with productivity parameters for each education level that vary between sectors. I assume worker preferences for firms depend on a worker's age, education level, and firm sector. Equipped with model predictions for matching and wages, I structurally estimate the model on matched employer-employee data. I estimate the model by maximum likelihood on the joint distribution of matching and wages, separately every three years.

The model developed in this paper is related both to one-to-many assignment problems studied in mechanism design (Bikhchandani and Ostroy (2002), Vohra (2011)), and to one-to-one matching models used in family economics (Choo and Siow (2006)). This paper bridges the gap between these two literatures: it extends one-sided assignments to two-sided matching, and generalizes one-to-one matching to one-to-many. Additionally, I extend the econometric framework of Choo and Siow (2006) and Galichon and Salanié (2021) to one-

to-many matching.

I use the novel theoretical framework developed to study the Portuguese labor market between 1987 and 2017. I highlight three facts on the Portuguese labor market: first, the country operates a vast education expansion over the period, which translates in a dramatic increase in the relative supply of high school graduates to non-graduates on the labor market. Second, the high school wage premium decreases over the period. The high school wage premium is defined as the wage gap between workers who graduated from high school, and those who did not. The decrease in wage premium is particularly stark among young workers. Third, I measure worker-firm sorting, which is defined as the relative number of high school graduates over non-graduates in an age group employed in a given sector. The distribution of high school graduates versus non-graduates across industry sectors becomes highly unbalanced, in favor of services, and transports and communications, who employ an increasing share of high school graduates. The former two facts imply relative supply of high school graduates over non-graduates has grown faster than firms' relative demand for high school graduates over non-graduates. The latter suggests that sorting between workers and firms has evolved over the period: either because firms in services and transport and communications demand an increasing share of high school graduates, or because high school graduates' preference for these firms strengthens.

Portugal is a particularly relevant example of rapid supply and demand changes on the labor market: it entered the European Union in 1986, which fuelled its economy's transition from being dominated by manufacturing (50% of the labor force employed in 1987), to services (30% of the labor force employed in 2017). Meanwhile, only 10% of its employed labor force held a high school degree in 1987, a percentage that has risen to 50% in 2017. As a point of comparison, the percentage of high school graduates in the US workforce has gone from 75% to 90% over the same period[1]. The proportional increase of high school graduates in Portugal is more extensive and starts from a much lower share of high school graduates on the labor market than in the US. In this respect, it is closer to the change in university graduates on the US labor market (from 20% to 35% over the same period). Graduating from high school has become much more common in Portugal over the last thirty years, but it is only in 2007 that high school graduates start representing the majority of young workers between 25 and 30. In 2017, 32% of the young workers between 25 and 30 still do not hold a high school degree. Meanwhile, university graduates in Portugal represented less than 3% of the employed labor force in 1987 and about 19% in 2017. Because the share of

[1] Percentages computed over workers aged more than 25, Census data

university graduates remains small for most of the period (it only reaches 10% in 2005), and because graduating from high school is still quite uncommon over most of the period I study, I consider a high school degree to be a differentiating signal in skill on the Portuguese labor market, much as a college degree is on the US labor market.

I find that relative demand for high school graduates from firms in the Services, Manufacturing, and Transport & Communications sectors has increased dramatically over the period, starting in the early 2010s. This finding is in line with the skill-biased technological change hypothesis. I also find that young and middle-aged high school graduates' preference for these industries has declined over time, while their share in production increases compared to senior workers. Compared to the classic supply and demand framework, these observations offer two additional mechanisms whereby high school wages gaps stay positive when a large number of high school-educated workers enter the labor market. First, a decrease in workers' amenities pressures wages upwards. Second, variation in young graduates' share in production compared to more senior high school graduates increases firm demands for the former compared to the latter. I perform several counterfactual exercises to assess the separate actions of changes in workers' demographics (both in education and age distribution), firm sector composition, firm demand through production parameters, and worker preferences, on sorting and wage premium. I find that changes in demographics are the main positive drivers of changes in sorting. Changes in industry composition, firm demand, and worker preferences overall have a negative, but modest, effect on sorting. Wage premia by age group and industry are negatively affected by changes in worker demography and industry composition and positively affected by changes in firms' demand. These suggest changes in relative productivity in favor of high school graduates have driven the high school wage premium up, but cannot compensate for the large increase in the relative supply of graduates versus non-graduates.

Related literature. The theoretical tools developed in this paper belong to the matching literature started by Becker (1973). My model is a one-to-many extension to the seminal work of Choo and Siow (2006) in the one-to-one case. As in Dupuy and Galichon (2022) and Galichon and Salanié (2021), it explicitly borrows tools from the optimal transport literature to introduce unobserved heterogeneity in the form of random utility and relies on Gretsky et al. (1992) to show equilibrium existence. This paper is also close to the hedonic model literature (Ekeland et al. (2004), Heckman et al. (2010)). A discussion of the links between hedonic models, matching with transferable utility, and optimal transport can be found in Chiappori et al. (2010). My work is also related to the seminal paper by Kelso

and Crawford (1982), and more recent work by Che et al. (2019) on one-to-many matching with non-transferable utility and Azevedo and Hatfield (2018) on one-to-many matching with transferable utility. They both show the existence of equilibrium for a large class of firm preferences, under a large market assumption, which I also use in this paper. I take one-to-many matching models a step further by taking the model described in this chapter to the data by introducing random shocks that account for unobservables and estimating it. The mechanism design literature has also explored many-to-one assignment problems in a one-sided framework with work by Bikhchandani and Ostroy (2002) and Vohra (2011).

This framework differs from the Sattinger model (Sattinger (1979), Sattinger (1993)) that assumes no unobserved heterogeneity and rests on the the firm's production function's supermodularity to find the optimal assignment of workers to firms. Fox (2010b) discusses non parametric identification of production functions in matching games and Fox et al. (2018) show that unobserved heterogeneity distribution can be recovered in matching games in which unmatched agents are observed and agents match on many separate markets. Because static random utility models (including mine) do not follow agents over time, they do not identify the unobserved heterogeneity distribution in the fashion of Abowd et al. (1999), Bonhomme et al. (2019), Bonhomme (2021) and instead focus on match formation based on observable surplus.

The model I develop features sorting between multidimensional types and as such is also related to Choné and Kramarz (2021), Lindenlaub (2017) and Lise and Postel-Vinay (2020). However, it is only remotely related to the search literature to which the latter paper belongs, as it focuses on relative supply and demand instead of search frictions. While the search literature often relies on Nash bargaining mechanisms, as in Shimer and Smith (2000) and Cahuc et al. (2006), the present model uses wage posting, as the competitive equilibrium in the model rests on wages that clear the labor market. Also related to this model and its application is the Roy model developed by Hsieh et al. (2019) to quantify the productivity gains of weakening discrimination barriers to women's and black men's entry into the labor market in the US. There exists an extensive literature on the education wage premium, mostly focused on the college wage premium in the US. Seminal work by Katz and Murphy (1992) shows that the increasing supply of college graduates in the 1970s and 1980s is absorbed on the US labor market by increased demand for these workers from firms. Card and Lemieux (2001) carry out a similar analysis that further differentiates workers by age, and show that young college graduates are the first to benefit from the slowdown in educational attainment in the 1980s. Goldin and Katz (2008) and Autor et al. (2020), among others, relate changes in the US wage structure to the race between education and technology, by which skill-biased technological change favors college graduates. Skill-biased

technological change (SBTC) origins in the development of new technologies, in particular computers (Autor et al. (1998), Autor et al. (2003)). However, if the SBTC hypothesis has proven a powerful explanation for the quick increase in graduate wage premium of the 1970s and 1980s, it is less clear if it can rationalize the subsequent slow down of both graduate wage premium and graduate supply in the 1990s, when the use of computers became prevalent (Card and DiNardo (2002)). The recent stagnation of the college wage premium in the US is also documented in several papers, and several explanations have been put forward: Beaudry et al. (2015) argue that the demand for cognitive skills has decreased since the early 2000s, pushing graduate workers down the job ladder. Valletta (2016) also emphasizes the role of job market polarization, i.e. the shift away from middle-skilled occupations, on college graduates' wages (as opposed to postgraduates). On the contrary, Blair and Deming (2020) examine job vacancy data and find that demand for skills has increased since the Great Recession. They explain the stagnating graduate wage premium by an increase in the supply of new graduates after 2008. They are backed by Hershbein and Kahn (2018) who show that the Great Recession has accelerated skill-biased technological change. In Portugal, changes in the wage structure are documented by Cardoso (2004), Centeno and Novo (2014) Almeida et al. (2017). To the best of my knowledge, I am the first to analyze the implications of worker and firm sorting on the education wage premium.

Outline. Section 2 describes the one-to-many matching model. Section 3 describes the evolution of the Portuguese high school wage premium between 1987 and 2017. Section 4 discusses the model's identification and estimation on Portuguese matched employer-employee data, and section 5 presents estimation results. Section 6 concludes.

2 Model

Recent administrative matched employer-employee datasets hold much more information than workers' characteristics and wage. They also inform on firms' characteristics and on matching, i.e. the joint distribution of workers and firms. Besides matching, the data also provides transfers between agents in the form of wage. Relying on this type of dataset enables to build a rich supply and demand framework to understand the race between education and technology. I build a one-to-many matching model where a single firm matches with several workers, who interact within the firm to produce output. Workers are compensated through wage, and hold specific preferences for different types of firms. Workers may also be unemployed. Firms maximize their profit, given their production function that is specific to their type and market clearing wage. Both worker and firm types are observed, and possibly

multidimensional. The model is an extension of Choo and Siow (2006) to a one-to-many framework, and existence of equilibrium rests on a large market assumption, as in Azevedo and Hatfield (2018) and Galichon and Salanié (2021). I model unobserved heterogeneity in the form of additive random utility. The social planner problem rewrites as a regularized optimal transport problem (Galichon (2016)), and I am therefore able to derive closed-form solutions for predicted matching and wage.

2.1 Setup

The labor market is two-sided, with workers and firms on each side. There is a continuum of workers $i \in I$. Each worker has a type $x \in \mathcal{X}$. Types are discrete and possibly multidimensional. There is a mass n_x of workers of type x, and a finite number of types: $\#\mathcal{X} = X$. On the other side of the market, there is a large number of firms $j \in J$. Each firm has a type $y \in \mathcal{Y}$. As for workers, firm types are also discrete and possibly multidimensional. There is a mass m_y of firms of type y, and a finite number of types: $\#\mathcal{Y} = Y$.

Each firm matches with a non-negative number of workers of each type, while each worker matches with a single firm. Let k_x be the number of type x workers a firm is matched with. The model is scaled by factor F, meaning that (n, m) and (Fn, Fm) are observationally equivalent. Hence the actual number of type x workers on the market is Fn_x. Therefore k_x must be comprised between 0 (a firm cannot hire a negative number of workers), and Fn_x. Vector k represents the workforce employed by the firm. It is akin to a bundle of workers of each type:

$$k = (k_1, \ldots, k_X) \in [0, Fn_1] \times [0, Fn_X].$$

Type x worker's utility for being employed at type y firm within workforce k is u_{xyk}. It is additive in a level of amenity α that depends both on worker and firm type, as well as workforce, and in wage w paid by the firm to the worker. Wage w_{xyk} is also allowed to depend on worker type, firm type and workforce.

$$u_{xyk} = \alpha_{xyk} + w_{xyk}.$$

Every worker also has the option to remain unemployed and obtain $u_{x0} = 0$.

Similarly, the firm profit v_{yk} is additive in production γ and minus total wage bill paid to its workforce.

$$v_{yk} = \gamma_{yk} - \sum_{x=1}^{X} k_x w_{xyk}.$$

Both amenity α_{xyk} and γ_{xyk} are functions of x, y, k and take their value in $\mathbb{R}$. The total surplus from a match between a firm and a workforce is the sum of workers' utilities and firm's profit

$$\Phi_{yk} = \sum_{x=1}^{X} k_x \alpha_{xyk} + \gamma_{yk}, \tag{2.1}$$

where wages have canceled out because they are modelled as perfectly transferable utility.

Some characteristics of firm and workers which play a role in match formation are unobserved, and therefore are not accounted for in x or y. There exists a large literature that deals with unobserved heterogeneity, and I build on a large subset (Choo and Siow (2006), Dupuy and Galichon (2014)) that uses additive random shocks to model it. I further assume a logit framework for the model by restraining the distribution of shocks to belong to the extreme value class, although as shown in Galichon and Salanié (2021) in the one-to-one case, identification is possible with a general class of distributions.

Worker i experiences stochastic shock $(\epsilon_{iyk})_{y,k}$ in addition to their systematic utility:

$$u_{x_i yk} + \xi \epsilon_{iyk}.$$

Similarly firm j experiences stochastic shock $(\eta_{jk})_k$ in addition to its systematic production:

$$v_{y_j k} + \xi \eta_{jk}.$$

where ξ is a scaling factor for unobserved heterogeneity. I impose the following independence conditions on stochastic shocks.

Assumption 2.1. *Stochastic shocks satisfy the following:*

(i) For each pair of two workers i and i', ϵ_{iyk} and $\epsilon_{i'yk}$ are mutually independent and identically distributed.

(ii) For each pair of two firms j and j', η_{jk} and $\eta_{j'k}$ are mutually independent and identically distributed.

(iii) For a worker i and a firm j, ϵ_{iyk} and η_{jk} are mutually independent.

(iv) ϵ_{iyk} is independent of $\alpha_{x_i yk}$, η_{jk} is independent of γ_{yk}.

(v) $(\epsilon_{iyk})_{y,k}$ and $(\eta_{jk})_k$ are distributed as extreme value 1 (Gumbel distribution).

A market is characterized by exogenous distributions of worker and firm types $(n_x)_{x\in\mathcal{X}}$ and $(m_y)_{y\in\mathcal{Y}}$, as well as amenity functions $(\alpha_{xy})_{x\in\mathcal{X},y\in\mathcal{Y}}$, production functions $(\gamma_y)_{y\in\mathcal{Y}}$, and

a draw of stochastic shocks ϵ and η. In the next subsection, I describe workers and firms choices and the resulting competitive equilibrium.

2.2 Competitive Equilibrium

Next, I define workers and firms expected utility and profit from choosing their best employer and/or workforce, given wages.

Definition 1. *Type x worker's expected indirect utility G_x as a function of u and type y firm's expected indirect utility H_y as a function of v are*

$$G_x(u_x) = \mathbb{E}\left[\max_{y,k}\left\{u_{xyk} + \xi\epsilon_{yk}, \xi\epsilon_0\right\}\right] \quad \textit{and} \quad H_y(v_y) = \mathbb{E}\left[\max_{k}\left\{v_{yk} + \xi\eta_k\right\}\right].$$

Under assumption 2.1, expected utilities rewrite in closed form.

Proposition 1. *Under assumption 2.1, expected indirect utilities write*

$$G_x(u_x) = \xi\log\left(1 + \sum_y\sum_k \exp\left(\frac{u_{xyk}}{\xi}\right)\right) \quad \textit{and} \quad H_y(v_y) = \xi\log\sum_k \exp\left(\frac{v_{yk}}{\xi}\right)$$

where $\sum_k = \sum_{k_1}\dots\sum_{k_X}$.

Proof. In Appendix C. □

The equilibrium on a market is found when supply from workers meets demand from firms. Supply and demand are defined as follows:

Definition 2. *Type x worker's supply is a vector $(S^x_{yk})_{yk,0}$ where S^x_{yk} is the mass of type x workers willing to match with type y firm and workforce k and S^x_0 is the mass of type x workers willing to remain unmatched.*

Type y firm's demand is a vector $(D^y_k)_k$ where D^y_k is the mass of type y firms willing to match with workforce k.

I model unemployment through S^x_0, which is determined at equilibrium. I assume no counterpart on the firm side: all firms must be matched to a given workforce.

Assumption on stochastic shocks lets us express supply from worker and demand from firms in logit form.

Proposition 2. *Under assumption* (2.1)*, the mass of type x workers willing to supply type y firms in workforce k is*

$$S_{yk}^{x} = n_x \frac{\exp\left(u_{xyk}\right)}{1 + \sum_{y,k} \exp\left(u_{xyk}\right)}. \tag{2.2}$$

The mass of type y firms who demand workforce k is

$$D_k^y = m_y \frac{\exp\left(v_{yk}\right)}{\sum_{k} \exp\left(v_{yk}\right)} \tag{2.3}$$

Proof. In Appendix C. □

Note that supply S and demand D both depend on wage schedule $w = \left(w_{xyk}\right)_{x,y,k}$. Because both workers and firms care not only about the other side's type, but also about the workforce they work with both in the systematic and stochastic parts of their utility or profit, wages also depend on workforce k. Therefore, two type x workers employed in two firms of same type y but who hire different workforce k and k' do not receive the same wage, as $w_{xyk} \neq w_{xyk'}$ in general. The model is able to generate heterogeneity in wage depending on firm size and workforce composition.

In the context of one-to-many matching, supply S and demand D are measured in different 'units': if a firm can match with several workers types, workers can only match with one firm type. Excess demand Z defined below gives the equivalence between worker and firm units.

Definition 3. *Given types x, y and workforce mass k, excess demand is defined as*

$$Z_{xyk}(w) = k_x D_k^y - S_{yk}^x.$$

A competitive equilibrium is reached on the market when supply and demand are feasible, matching is incentive compatible, and excess demand is zero. The first two conditions are automatically filled as a byproduct of the definition of supply and demand: in proposition 2, workers and firms choose their optimal option. As a result, matching is incentive compatible, and supply and demand are feasible:

$$\sum_{y,k} S_{yk}^x + S_0^x = n_x \quad \text{and} \quad \sum D_k^y = m_y.$$

Definition 4. *An equilibrium outcome (S, D, w) satisfies $\forall x, y, k$: $Z_{xyk}(w) = 0$.*

The existence a competitive equilibrium rests on the fact that there are large numbers of agents on the market. To show existence, I follow a proof technique introduced in the continuum assignment problem by Gretsky et al. (1992), and already used for one-to-one matching markets by Galichon and Salanié (2021). The reasoning is also very close to Azevedo and Hatfield (2018)'s proof for competitive equilibrium existence in a large economy on a market of buyers and sellers with a finite set of possible trades. Bikhchandani and Ostroy (2002) explore a similar assignment problem but do not assume large markets and work without heterogeneous shocks.

I prove existence of equilibrium in two steps. First, I show that the competitive equilibrium reframes as an optimization problem on total welfare. Second, I show that this problem is the dual of the social planner problem, who maximizes total surplus under feasibility conditions. The social planner problem maximizes a continuous and strictly concave function over a compact space. As such, a unique solution exists.

Theorem 2.1. *Equilibrium payoffs obtain as solutions to the following problem:*

$$\begin{aligned} &\inf_{u,v} \sum_x n_x G_x(u_x) + \sum_y m_y H_y(v_y) \\ &s.t \sum_x k_x u_{xyk} + v_{yk} = \Phi_{yk} \quad \forall k, y. \end{aligned} \tag{2.4}$$

Proof. In Appendix C □

Theorem 2.2. *Equilibrium matching* $\mu_{yk} = D_k^y = \frac{S_{yk}^x}{k_x}$[2] $\forall x$ *and equilibrium* S_0^x *obtain as solution to the social planner problem:*

$$\begin{aligned} &\max_{\mu, S_0} \sum_y \sum_k \Phi_{yk}\mu_{yk} + \xi\mathcal{E}(\mu, n, m) \\ &s.t \sum_y \sum_k k_x \mu_{yk} + S_0^x = n_x \\ &\sum_k \mu_{yk} = m_y, \end{aligned} \tag{2.5}$$

[2] Equality $\mu_{yk} = \frac{S_{yk}^x}{k_x}$ is only defined when $k_x > 0$. If $k_x = 0$, supply S_{yk}^x is not defined

where $\mathcal{E}(\mu, n, m)$ is equal to

$$\mathcal{E}(\mu, n, m) = -\sum_x n_x \sum_y \sum_k \frac{k_x \mu_{yk}}{n_x} \log \frac{k_x \mu_{yk}}{n_x} - \sum_x n_x \frac{S_0^x}{n_x} \log \frac{S_0^x}{n_x}$$
$$- \sum_y m_y \sum_k \frac{\mu_{yk}}{m_y} \log \frac{\mu_{yk}}{m_y}.$$

The solution to (2.5) *exists and is unique.*

Proof. In Appendix C. □

Theorem 2.2 shows that equilibrium matching can be obtained by solving a penalized social planner problem, where the objective function is the difference between total expected surplus and an entropy term due to unobserved heterogeneity. It is reminiscent of the discrete regularized optimal transport problem (Galichon (2016)). However it differs from the usual transport problem in two important ways: first workers are allowed to remain unmatched through S_0^x, and second, the first marginal condition $\sum_y \sum_k k_x \mu_{yk} + S_0^x = n_x$ is not a condition on the marginal distribution of k, which is endogeneous, but on the marginal distribution of worker types.

Solving for problem (2.2) yields the following expressions for equilibrium matching μ, unemployment S_0^x and wages w.

Proposition 3. *Equilibrium matching solves*

$$\begin{aligned} \log \mu_{yk} &= \frac{\Phi_{yk} - \sum_x k_x U_x - V_y + \xi \sum_x k_x \log \frac{n_x}{k_x} + \xi \log m_y}{\xi(1 + \sum_x k_x)} \\ \log S_0^x &= \frac{-U_x + \log n_x}{\xi}. \end{aligned} \tag{2.6}$$

Equilibrium wages write

$$\begin{aligned} w_{xyk} =& \frac{\gamma_{yk} - \alpha_{xyk} + U_x - V_y + \xi \log m_y - \xi \log \frac{n_x}{k_x}}{\xi(1 + \sum_x k_x)} \\ &+ \frac{\sum_{x' \neq x} k_{x'} \left((\alpha_{x'yk} - \alpha_{xyk}) - (U_{x'} - U_x) + \xi \log \frac{n_{x'} k_x}{n_x k_{x'}} \right)}{\xi(1 + \sum_x k_x)}. \end{aligned} \tag{2.7}$$

Where U_x, V_y solve

$$\begin{cases} \sum_{y,k} k_x \exp\left(\frac{\Phi_{yk} - \sum_x k_x U_x - V_y + \xi \sum_x k_x \log\left(\frac{k_x}{n_x}\right) + \xi \log m_y}{\xi(1+\sum_x k_x)}\right) + \exp\left(\frac{-U_x + \xi \log n_x}{\xi}\right) = n_x \\ \sum_k \exp\left(\frac{\Phi_{yk} - \sum_x k_x U_x - V_y + \xi \sum_x k_x \log\left(\frac{k_x}{n_x}\right) + \xi \log m_y}{\xi(1+\sum_x k_x)}\right) = m_y. \end{cases} \tag{2.8}$$

Proof. In Appendix C □

In practise, equilibrium μ, S_0^x and w are computed by solving for equations (2.8) using the Sinkhorn algorithm, also called IPFP (Iterative Proportional Fitting Procedure), that has been developed in the optimal transport literature (among others). In the one-to-many case, U_x and V_y can be solved for by coordinate update in the same spirit as Sinkhorn.

2.3 Links with search and matching models in the literaure

The model I develop is akin to Choo and Siow (2006)'s in a one-to-many instead of a one-to-one setting. One can view the space of workforces, instead of workers, as a side of the market, with firms on the other side. It is particularly striking that just like in Choo and Siow (2006), both equilibrium matching and wage are weighted by the number of individuals in the match $1 + \sum_x k_x$. In this representation, the model almost reduces to the one-to-one framework, but for the specific shape of marginal conditions in (2.8), that links the matching over workforces and firms back to the number of workers of each type. Another difference with Choo and Siow (2006), Dupuy and Galichon (2022) and other frameworks that use the IPFP algorithm in their framework is that expected indirect surpluses U and V cannot be explicitly expressed through equations (2.8) because the size of every match is endogenous. I observe transfers as wages and can leverage them to split total match surplus between workers and firms, in the spirit of Dupuy and Galichon (2022).

The model also features wage posting. In the decentralized equilibrium, firms choose among workforces and associated wages given their draw of random shock η, while workers choose among firm types, workforces and wages given their draw of ϵ. A salient feature of the model is that it generates wage dispersion for a given worker and firm type, based on the workforce hired by the firm. All other things equal, wage is increasing in the number of workers hired by the firm. This is reminiscent of search models such as Burdett and Mortensen (1998), although the model presented here is not a search model.

Finally, my model is closer to Katz and Murphy (1992) and Card and Lemieux (2001)

than it may appear at first sight. To see this, consider two workforces k and k', where $k'_x = k_x$, expect for $k'_{\bar{x}} = k_{\bar{x}} + t$, i.e. there is t more worker of type $\bar{x}$ hired in workforce k'. Then firm production and type $\bar{x}$ worker's wage satisfy:

$$\gamma_{yk} - \gamma_{yk'} = \left(1 + \sum_x k_x\right) w_{\bar{x}yk} - \left(1 + \sum_x k'_x\right) w_{\bar{x}yk'}.$$

At the limit, when t tends to zero (if the extra worker works very few hours for instance), we obtain the same intuition as with the representative firm that the marginal change in wage is equal to the marginal change in production (divided by the number of agents):

$$\frac{\partial \gamma_{yk}}{\partial k_x} = \left(1 + \sum_x k_{\bar{x}}\right) \frac{\partial w_{\bar{x}yk}}{\partial k_{\bar{x}}}.$$

Hence any change in workers' $\bar{x}$ is proportional to their marginal productivity, although its impact is mitigated by total number of workers hired by the firm.

3 Empirical Evidence

3.1 Data Description

The *Quadros de Pessoal* dataset offers an exhaustive snapshot of the Portuguese labor market every year from 1987 to 2017. It covers all employees in the private sector (except domestic workers), and provides information on their age and highest degree obtained, as well as their monthly wage and hours worked. To compute the high school wage premium by age, I part the worker population into two groups: those who did not graduate from high school, and those who did. I also categorize workers into three age groups: young workers (from 16 to 35 years old), middle aged workers (from 36 to 50 years old), and senior workers (from 51 to 68 years old). I only consider full time employees, that is, workers that are neither part time workers (approximately 10% of the observations) nor self-employed, in unpaid family care, or in other forms of employment (less than 1% of the observations). I compute real hourly wage as the ratio of monthly wage over monthly hours, controlling for inflation and clean out the lowest 1% and highest 99% hourly wage percentiles. Firms belong to either five sectors, or industries: primary industries (agriculture, mining, energy, construction), manufacturing, retail and hospitality, services, transport and communications.

To account for unemployment, I use public yearly unemployment figures by education

level and age group provided by INE[3]. Information on unemployment is missing between 1987 and 1991, hence I assume the unemployment rate in these years is the same as in 1992. I compute the number of unemployed workers each year by education level and age group by combining unemployment rates and the number of observed employed workers in *Quadros de Pessoal*. In what follows, active worker refers to workers either employed or unemployed.

3.2 Empirical facts

The Portuguese labor market is characterized by three facts between 1987 and 1997. The first is the dramatic increase in the number of high school educated workers, compared to the number of workers who did not go to high school. The second is the decrease in high school wage premium, i.e. the wage gap between high school graduates and non graduates. The third is the change in sorting between education level on the worker side, and industry on the firm side: sorting intensity between high school graduates and specific industries rises over the period. Each of these three facts are detailed below.

Fact 1 : Education supply. Supply of high school graduates relative to non-graduates rises dramatically over the period, as evidenced by Figure 2.1. Relative supply is measured as the ratio of number of high school graduates over number of active school graduates by age group in each year. Because high school enrolment grows every year, young workers are more impacted by this growth, and their relative supply goes from .12 to 1.79 on Figure 2.1, meaning high school graduates have grown to be about eight times less numerous to almost twice as numerous as non-graduates between 1987 and 2017.

[3]Found on their website

Figure 2.1: High school graduates versus non graduates relative supply, by age group

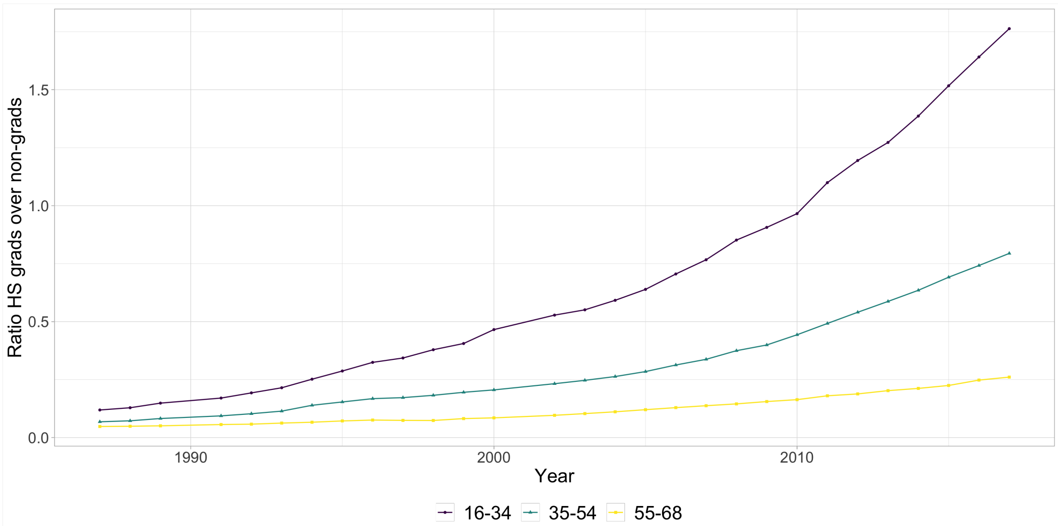

Fact 2 : Wage premia by age group. The second fact that characterizes the Portuguese labor market is the decrease in high school wage premium. To compute high school wage premium by age group, I estimate the following equation by OLS:

$$\log w_{ijt} = \sum_{a_i \in \{y,m,s\}} \mathbb{1}_{[\text{HS graduate}_i]} \beta_{a_i t} + g_i + r_{jt} + d_{jt} + u_{ijt}, \tag{2.9}$$

where each individual i working in firm j at time t earns wage w_{ijt}. a_i is individual i's age group: either y, m or s. $\mathbb{1}_{[\text{HS graduate}_i]}$ equals 1 if i graduated from high school, and 0 otherwise. g_i, r_{jt} and d_{jt} are gender, region and industry fixed effects. β_{at} is the yearly high school wage premium, differentiated by age group: it measures how much more in percentage a high school graduate earns compared to a non high school graduate. I allow fixed effects to vary over time, I estimate (2.9) separately every year.

Figure 2.2 shows the change in estimated high school wage premium over time for each age group, along with 95% confidence intervals. The high school wage premium differs between age groups: the wage gap is much higher (between 60% and 80% over the period) for senior workers than for younger workers (between 40% and 20%). Figure 2.2 also shows that the wage premium decreases for all age groups between 1987 and 2017. The extent of the decrease is different depending on age however: senior workers lose only about 17 percentage points (p.p) in high school wage premium over the period, while young workers lose almost 50p.p and middle ages workers lose slightly less than 30p.p.

Figure 2.2: Estimated high school graduate wage premium by age group

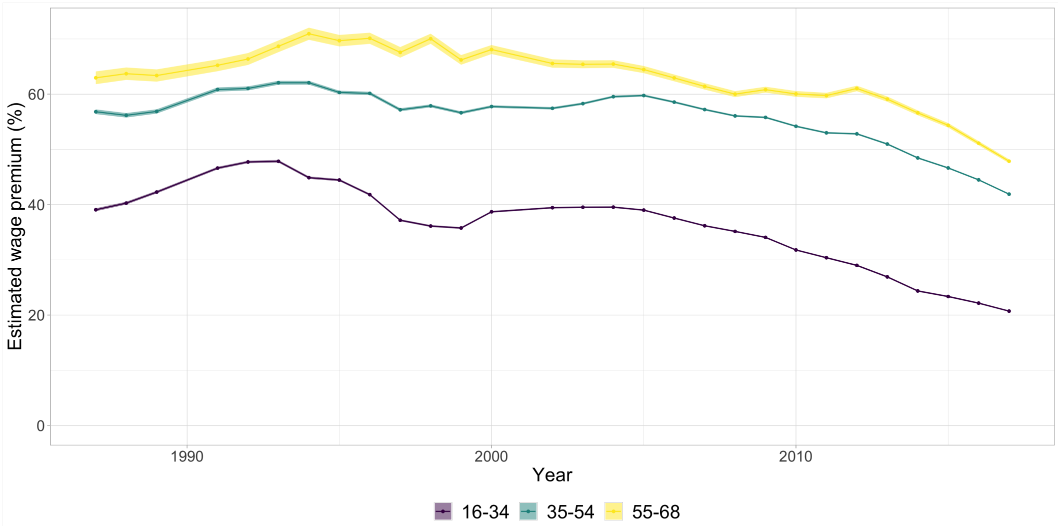

Wage levels differ by gender, with men earning more on average than women in all education levels and age groups (see Figures 2.9 and 2.10 Appendix B). Yet, both men and women experience the trend described in Figure 2.2: the high school wage premium decreases for both genders, more strikingly for young men and women.

Unlike most of the literature, Fact 2 focuses on the high school wage premium, rather than the college wage premium. This choice stems from the particular set up of Portugal: on the Portuguese labor market, a high school degree is a defining factor in a workers career, because it is less common than in other developed economies. For instance in 2017, 32% of young workers still do not hold a high school degree. However, it is important to note that the university wage premium, defined as the wage gap between university graduates and non graduates, follows a similar trend to the high school wage premium, as shown in Appendix B, Figure 2.11.

Fact 3 : Sorting between education levels and sector. Sorting between education level and industry is measured by age group as the ratio of the number of employed high school graduates to employed non-graduates in a sector. Sorting is stronger between high school graduates and sector A than sector B, if this ratio is larger in sector A than in sector B. Plotting sorting ratios by sector over time reveals stark differences by industry, as shown in figure 2.3. Most notably, the Services and Transport and Communications industries hire young high school graduates over non-graduates at a higher rate than the change in overall relative supply. As shown in fact 1, relative supply goes from .11 to 1.79 over the period, while the sorting ratio in these industries reaches 3.22 and 4.34 in 2017. Services and

transports and communications also hire proportionally more middle-aged workers, with a ratio of 1.61 and 1.39 in 2017, compared to a relative supply ratio of .82.

Figure 2.3: High school - Sector sorting, by age group

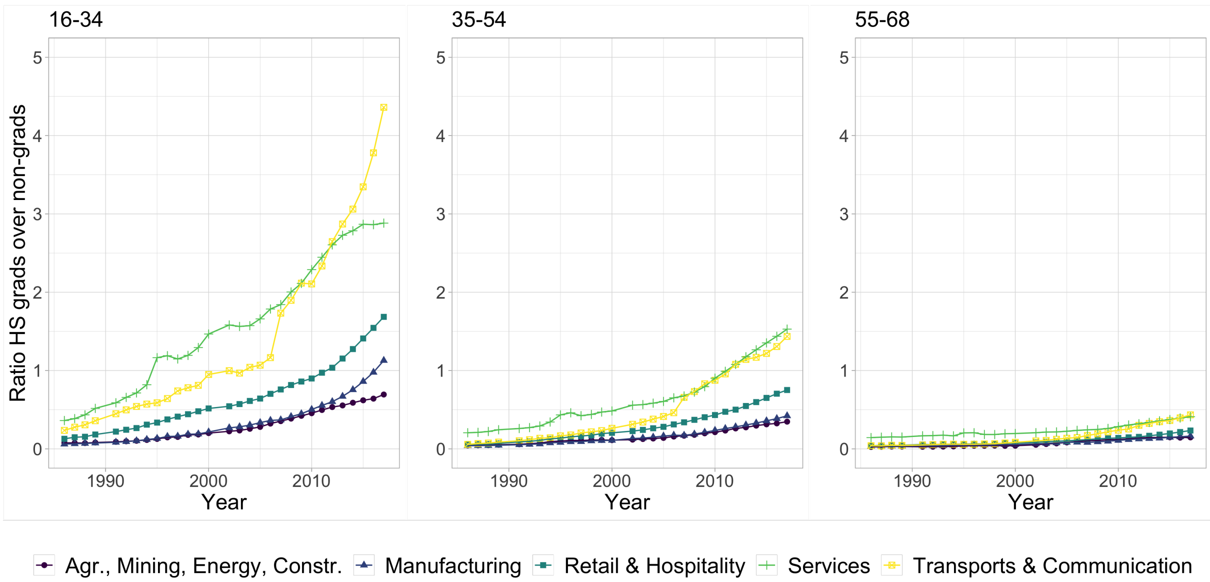

Summary. The relative supply of high school graduates over non graduates rises for all age groups, and in particular among young workers. Meanwhile, the high school wage premium decreases in Portugal between 1987 and 2017. Its decline is particularly strong for young workers, between 16 and 34 years old. The rise in relative supply is not absorbed equally by all sectors: Services and Transports and Communications hire proportionally more young and middle-aged high school graduates than other sectors. This is indicative of strong sorting between these workers and the Services and Transports and Communications industry.

Portugal is unique in that it has known a dramatic education expansion, going from 10% of high school graduates in the labor force in 1987 to about 50% in 2017. It has also known deep changes in how workers sort with firms based on education level, age group, and the firm sector, as evidenced in Fact 2. As such, it is an ideal laboratory to understand how sorting between workers and firms drives the high school wage premium over time. Changes in sorting can be caused either by an increase in relative productivity of high school graduates in some industries, a change in preferences of young high graduates, or changes in substitution patterns among education levels or age groups. Meanwhile, the increase in relative supply of high school graduates likely drives the wage premium down. The wider economic interpretation could go in two different ways: a trade effect or a technology effect. Indeed, Portugal entered the European Union in 1986, which lowered barriers to

trade with other EU countries. Because Portugal has relatively more uneducated workers than the countries it trades with, a Heckscher-Olhin model of trade predicts it will start exporting more goods which require uneducated labor, which would increase firms' demand for uneducated workers. The technology effect postulates that, like other Western countries over the period, Portugal has experienced skill-biased technological change, which instead would increase demand for educated workers. The decreasing wage premium would then be explained by the relative increase in educated labor supply, which outbalances the rise in demand. To tell apart these two interpretations, I parametrize in the next section the model presented in section 2 to untangle the effect of changes in relative supply from changes in firm production and worker preferences, and evaluate their impact on sorting and wage.

4 Identification and Estimation

Rearranging equations (2.6) and (2.7) (see appendix C), we obtain that amenities α and production γ verify the following equations:

$$\begin{aligned} \alpha_{xyk} &= U_x - w_{xyk} + \xi \log \mu_{yk} - \log \frac{n_x}{k_x} \\ \gamma_{yk} &= V_y + w_{xyk} + \xi \log \mu_{yk} - \log m_y. \end{aligned} \tag{2.10}$$

Hence α_{xyk} and γ_{yk} are identified up functions U_x and V_y. Inspection of (2.6) and (2.7) shows a model generated by $\alpha_{xyk} + a_x$ and $U_x - a_x$ for any a_x, would be observationally (i.e. matching and wage would be the same) equivalent to a model generated by α_{xyk} and U_x, if it was not for the fact that unemployed workers are accounted for, and that their amenities are assumed to be 0. Because U_x and $U_x + a_x$ do not generate the same share of unemployed workers, α_{xyk} is identified from observing the share of unemployed workers. Single firms (who would not employ anyone) are not observed however, so that two models generated by γ_{yk} and V_y or $\gamma_{yk} + b_y$ and $V_y - b_y$ are observationally equivalent. Therefore, I set any variable that varies in firm type y but is constant across workforce k to zero in firm production by assuming a Constant Elasticity of Substitution in the parametrization.

The model's predictions on matching (2.6) and wage (2.7) allow to separately identify amenity and productivity functions $(\alpha_{xy})_{xy}$ and $(\gamma_y)_y$. This would not be true if we observed only matching, as α and γ appear together in the matching prediction, and only total surplus can be identified from this equation. If only wages were observed, the same problem arises and only the difference between firm production and worker amenities is identified. In this case one must assume that amenities are zero in order to identify production.

Note that the model identification does not rest on observation of demand or supply shifters. Instead, it uses variation in matches within agents of the same type. As a though experiment, imagine there would only be one type of workers. Then under the assumptions made on unobserved heterogeneity, the number of workers matched with a type y firm relative to how many are matched with a type y' firm informs the model on the exact amenity difference perceived by workers between firms of type y and firms of type y'. Because I assumed that unmatched workers perceive no amenity, these differences translate into amounts of amenities, that are always relative to the unmatched worker's, as is usual in this type of logit models.

In any given period t, I aim at parametrically estimate α^t and γ^t. All amenity and production parameters are allowed to vary with time, and in what follows I drop the superscript t to ease the exposition. I assume $N = 6$ worker types that are the combination of two education levels, and three age groups. The education levels are high school graduates H and non graduates L, and the age groups are young y (below 35), middle-aged m (between 35 and 54), and senior s (above 55). Let $e(x)$, $a(x)$ be type x's education level and age group. Firm workforce k is composed of the numbers of each worker type employed

$$k = (k_{H,y}, k_{H,m}, k_{H,s}, k_{L,y}, k_{L,m}, k_{L,s}) \,.$$

Employed number of worker k_x is directly observed in the data and defined as total number of hours worked monthly by workers of type x hired by the firm, divided by 174, the monthly hours equivalent of a 40 hours week. Hence each k_x counts the full-time equivalent of the number of type x workers employed by the firm. This measure is not necessarily an integer, as part-time workers would count as fractions of the full-time equivalent. Type y firm produces according to a nested Constant Elasticity of Substitution (CES) production function with different parameters depending on its type y:

$$\gamma_{yk} = \left[(\theta_H^y H(k))^{\frac{\sigma-1}{\sigma}} + (\theta_L^y L(k))^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} - \sum_x \mathbb{1}_{[k_x>0]} \frac{\nu^y}{n_x},$$

where aggregates $H(t)$ and $L(t)$ are:

$$H(k) = \left[\sum_{a\in\{y,m,s\}} \lambda_{a,H} k_{H,a}^{\frac{\tau^H-1}{\tau^H}} \right]^{\frac{\tau^H}{\tau^H-1}} \quad \text{and} \quad L(k) = \left[\sum_{a\in\{y,m,s\}} \lambda_{a,L} k_{L,a}^{\frac{\tau^L-1}{\tau^L}} \right]^{\frac{\tau^L}{\tau^L-1}} .$$

Production γ^y's outer nest involves three parameters: σ, θ^y_H, θ^y_L and two aggregate inputs $H(k)$ and $L(k)$. $\sigma \in (0, \infty)$ is the elasticity of substitution between education levels, it is greater than one if high school graduates and non graduates are gross substitutes, and smaller than one if they are gross complements. σ is assumed to be the same across firm types. θ^y_H, $\theta^y_L \in [0, \infty)$ are graduates and non graduate's productivity parameters. Both parameters may vary by firm type y. In addition to their CES production function, firms experience friction $\frac{\nu^y}{n_x}$ if they employ workers of type x. The rationale is that if there are few workers of type x, then it is costly for the firm to find and hire them. ν^y measures this cost by sector y.

The CES parametrization is the last piece needed for complete identification of production γ: indeed, there does not exist a constant b_y such that γ_{ky} and $\gamma_{yk}+b_y$ are observationally equivalent across all workforces k under the CES assumption. This is true up to search friction $\sum_x \mathbb{1}_{[k_x>0]}\frac{\nu^y}{n_x}$, hence search costs ν^y are not comparable across sectors.

Aggregate labor inputs $H(k)$ and $L(k)$ form the production function's inner nest. They each depend on four parameters: three age productivity parameters each: $\lambda^{y,t}_{a,H}$ and $\lambda^{y,t}_{a,L} \in [0, \infty)$ and one elasticity of substitution between age levels each: τ^H and $\tau^L \in (0, \infty)$. Elasticities vary by education level but are the same across firm types, while age productivity vary with firm type y.

The production function is close to the one used by Katz and Murphy (1992), and Card and Lemieux (2001): it assumes imperfect substitution and varying productivity in the tasks performed by different education levels and age categories. Capital is not included as an input, but may impact productivity parameters through firm type: if two firm types use different levels of capital in relation to education levels, it is reflected in the levels of θ^y_H and θ^y_L. Unbiased technological change that increases all workers productivity results in an increase in both θ^y_H and θ^y_L. Technological change may be biased towards an education level if its productivity increases faster than the other's. This production function also allows more flexibility than Card and Lemieux (2001) by letting elasticities of substitution and age productivity vary in time.

Production assumes constant returns to scale. Note that it is homogeneous of degree one, and therefore two functions parametrized with θ and λ or $c \times \theta$ and $\frac{\lambda}{c}$ are equivalent. To distinguish between these versions, I impose normalization condition:

$$\sum_a \lambda^y_{a,H} = \sum_a \lambda^y_{a,L} = 1 \quad \forall y. \tag{2.11}$$

I assume worker amenities are constant in k:

$$\alpha_{xyk} = \beta_x^y. \tag{2.12}$$

β_x^y reflects type x worker preferences for type y firms over other firm types. In particular I assume workers are indifferent to workforce size.

Given these functional forms, I am looking to estimate in every period t parameters $\left(\lambda_{a,H}^y\right)_a$, $\left(\lambda_{a,L}^y\right)_a$, $(\theta_H^y)_y$, $(\theta_L^y)_y$, $(\beta_x^y)_{x,y}$, τ_H, τ_L and σ. To this aim I use a maximum likelihood method, which I describe in what follows.

The model predicts matching μ_{yk} as a joint distribution on firms and workforces, which can be compared to observed matching $\tilde{\mu}_{yk}$, which is simply the number of firms matched with workforces k in the data. Let also $\tilde{S}_0^x$ be the number of unemployed worker of type x. Let $\tilde{w}_{ij}$ be the observed wage of worker i employed by firm j. Observed wage $\tilde{w}_{ij}$ is assumed to be a noisy measure of predicted wage $w_{x_i y_j k_j}$ where k_j is the entire workforce employed at firm j. In other words:

$$\tilde{w}_{ij} = w_{x_i y_j k_j} + \nu_{ij} \text{ where } \upsilon_{ij} \sim \mathcal{N}(0, s^2) \text{ iid}, \tag{2.13}$$

where υ_{ij} is a centered measurement error of variance s^2. Under assumption (2.13), observed average wage $\tilde{W}_{xyk}$ for type x workers hired by firm y in workforce k is distributed as

$$\tilde{W}_{xyk} = \frac{1}{\tilde{K}_{xyk}} \sum_{\substack{i:x_i=x \\ j:y_j=y}} w_{x_i y_j k_j} \sim \mathcal{N}\left(0, \frac{s^2}{\tilde{K}_{xyk}}\right) \text{ iid}, \tag{2.14}$$

where $\tilde{K}_{xyk}$ is the total number of type x workers hired by firm y in workforce k in the data: $\tilde{K}_{xyk} = k_x \tilde{\mu}_{yk}$. Because there is a very large number of observed wages in the data (as many as there are workers), I choose to work with observed average wages by worker type, firm type and workforce in the likelihood estimation. This reduces the likelihood function complexity but does not limit estimation: the model parameters as well as variance s^2 can still be recovered from log likelihood maximization.

Let $\mu_{yk}(\Gamma, \beta, n, m)$ and $w_{xyk}(\Gamma, \beta, n, m)$ be the matching and wage predicted by the model, given parameters $\Gamma = ((\theta_H^y)_y, (\theta_L^y)_y, (\lambda_{H,a})_a, (\lambda_{L,a})_a, \tau_H, \tau_L, \sigma)$, β, and worker and firm type

distributions $n = (n_x)_x$, $m = (m_y)_y$. The log likelihood of observing pair $(x, y, k, \tilde{W})$ is then

$$\begin{aligned} &k_x \tilde{\mu}_{yk} \log \mu_{yk}(\Gamma, \beta, n, m) \\ &\quad - \tilde{K}_{xyk} \frac{(\tilde{W}_{xyk} - w_{xyk}(\Gamma, \beta, n, m))^2}{2s^2} - \frac{1}{2} \log \left(\frac{s^2}{\tilde{K}_{xyk}} \right). \end{aligned}$$

Meanwhile, the log likelihood of observing an unemployed worker of type x is

$$\tilde{S}_0^x \log S_0^x(\Gamma, \beta, n, m, s^2).$$

The log likelihood method therefore solves

$$\begin{aligned} &\max_{\Gamma,\beta,s^2} l(\Gamma, \beta, n, m, s^2) \\ &= \max_{\Gamma,\beta,s^2} \sum_x \sum_{y,k} k_x \tilde{\mu}_{yk} \log \mu_{yk}(\Gamma, \beta, n, m, s^2) + \sum_x \tilde{S}_0^x \log S_0^x(\Gamma, \beta, n, m, s^2) \\ &\quad - \sum_x \sum_{y,k} \tilde{K}_{xyk} \frac{(\tilde{W}_{xyk} - w_{xyk}(\Gamma, \beta, n, m, s^2))^2}{2s^2} - \frac{1}{2} \log \left(\frac{s^2}{\tilde{K}_{xyk}} \right). \end{aligned} \tag{2.15}$$

I run log likelihood estimation on ten separate three year periods between 1987 and 2017[4]. Years in each period are pooled. In each period, I observe number of workers and firms $(\tilde{n}_x)_x$ and $(\tilde{m}_y)_y$ directly in the data. I normalize without loss of generality the total mass of firms in each period to 1, so that scaling factor F is $\sum_y \tilde{m}_y$, and input $n_x = \frac{\tilde{n}_x}{F}$ and $m_y = \frac{\tilde{m}_y}{F}$ to likelihood estimation.

I solve numerically for problem (2.15) using a nested method: in the inner loop, $\mu(\theta, \lambda, \tau, \sigma, \beta)$, $S_0^x(\theta, \lambda, \tau, \sigma, \beta)$ and $w_{xyk}(\theta, \lambda, \tau, \sigma, \beta)$ are computed according to (2.6) and (2.7). Scaling factor ξ is set to 1. In the outer loop, I update $(\theta, \lambda, \tau, \sigma, \beta)$ using Adam, a gradient descent method with momentum (Goodfellow et al. (2016), Kingma and Ba (2017)). Variance s^2 is obtained in the outer loop through first order condition:

$$s^2 = \frac{1}{W} \sum_x \sum_{y,k} \tilde{K}_{xyk} \left(\tilde{W}_{xyk} - w_{xyk}(\Gamma, \beta, n, m) \right)^2.$$

More details on estimation can be found in appendix D.

[4]Periods are 1987-1989, 1991-1993, 1994-1996, 1997-1999, 2000-2003, 2004-2006, 2007-2009, 2010-2012, 2013-2015, 2016-2017. Since data for years 1990 and 2001 are missing, the last time period spans only two years.

5 Results

5.1 Parameters estimates

Estimates for high school graduates and non-graduates productivities θ_H^y and θ_L^y by industry y are presented in figure 2.4. Estimates shows education productivity are heterogeneous by industries, and have evolved in non-linearly: high school graduates productivity displays an impressive surge starting in 2010, especially in the Transport & Communications, Manufacturing and Services industries. Non-graduates productivity drops for all industries between 2010 and 2013. As a result, high school graduates' productivity relative to non-graduates rises at the end of the period.

Figure 2.4: Estimated education productivities

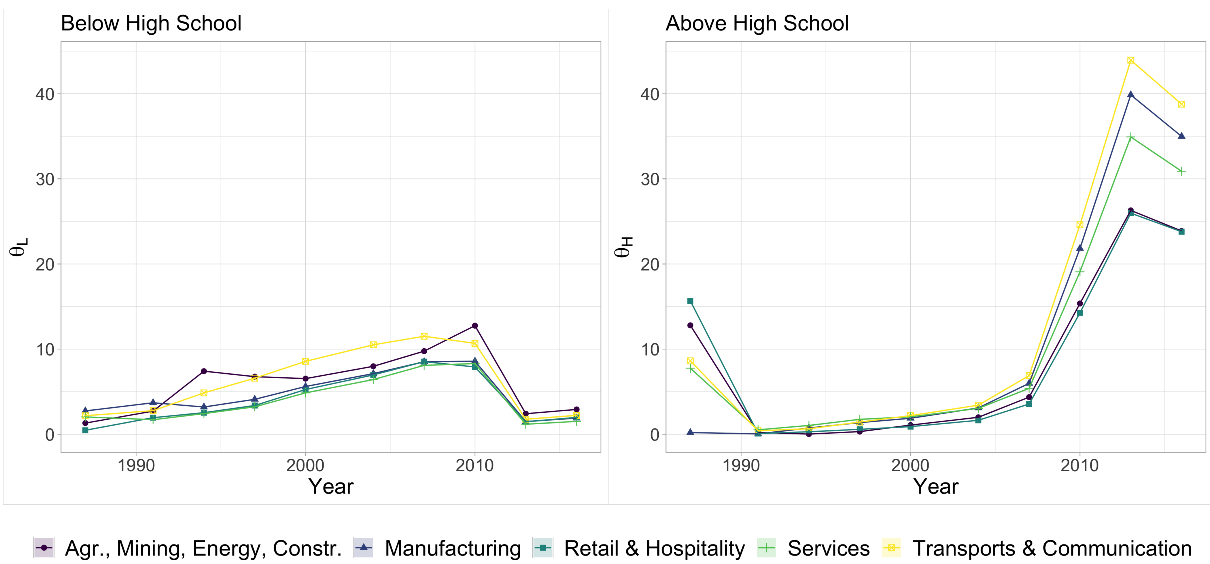

95% confidence interval

The spectacular increase in educated workers' productivity that starts between 2007 and 2010 and stabilizes after 2013 coincides with the sovereign debt crisis in Portugal, which may give some clues as to why the increase is so large. Indeed, in an attempt to curb unemployment and stimulate the labor market, Portuguese labor institutions have been modified on several levels: minimum wage, which had steadily increased in the previous years, was frozen, severance pay was lowered, and an attempt was made to revise the bargaining of wages at the sectoral level, although this attempt met with a lot of resistance and never entirely went through. The first two changes might have had an impact on worker productivity as measured by the model on their own however: because estimated productivity is positively tied to wage, the yearly increases in minimum wage have driven the increase in uneducated

workers productivity before 2010. When minimum wage froze, this driver of growth did too. Second, the lowering of severance pay seems to have accelerated the replacement of uneducated workers by educated workers in firms: in 2007, 48.8% of firms employed at least one educated worker. In 2010, this percentage was up to 53.0% and in 2013 to 57.7%. Hence, the matching distribution has drastically changed in these years, which is likely driving the rise in educated workers' productivity, and contraction in uneducated workers' productivity. Under this interpretation, the sudden changes in trend observed in 2010 reflect an overdue adjustment of matching on the labor market, made possible by the change in labor market institution.

Figure 2.5 shows the evolution of age productivities $\lambda_{H,a}$ and $\lambda_{L,a}$ by age group and over time. Because $\lambda_{H,a}$ and $\lambda_{L,a}$ sum to one over age groups in any given year, they can be interpreted as shares of each age group in total labor input by high school graduates and non graduates. Estimates $\lambda_{L,a}$ are fairly stable over the period up until 2010-2013, with middle-aged workers making up most of the labor input for non high school graduates. Their share in labor input increases to an even higher level (about 75%) in 2010-2013. Estimates $\lambda_{H,y}$ and $\lambda_{H,m}$ increase steadily until the early 2000s, but high school graduates senior workers input remains the most productive of the three at the end of the period.

Figure 2.5: Estimated age productivities

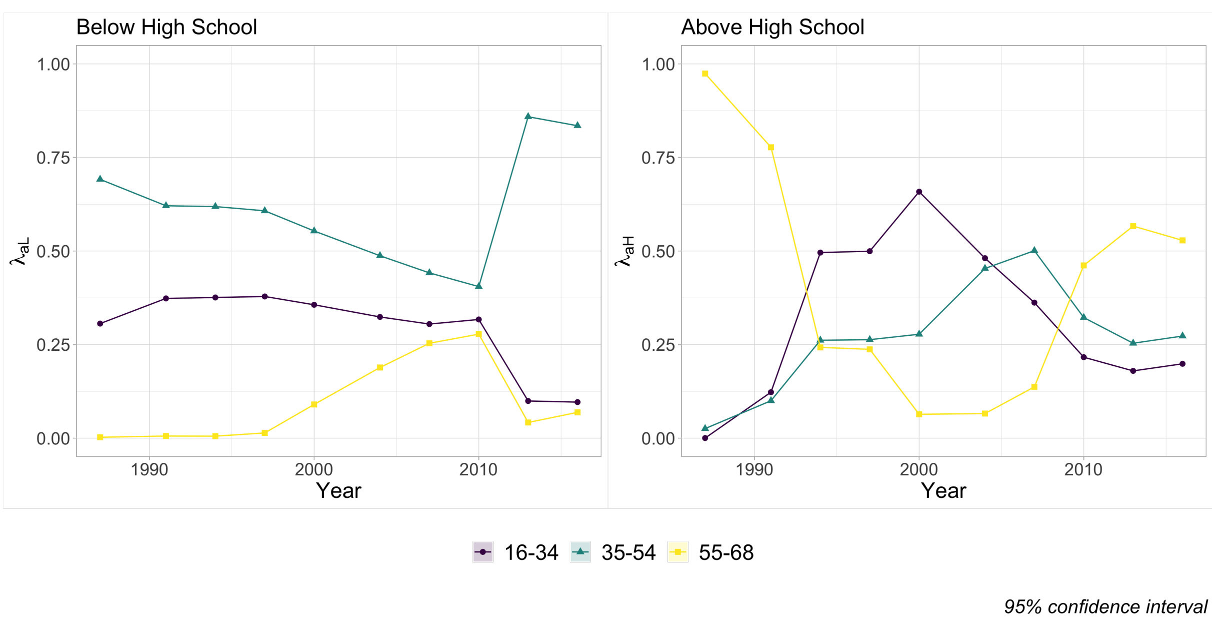

Figure 2.6 presents the change in worker preferences for firms β_x^y in euros per hours worked. All education levels and age groups hold high preferences for Retail & Hospitality over the period, and low preference for Transport & Communications. High school graduates'

preference for Services increases over the period, while their preference for Manufacturing decreases.

Figure 2.6: Estimated worker preferences

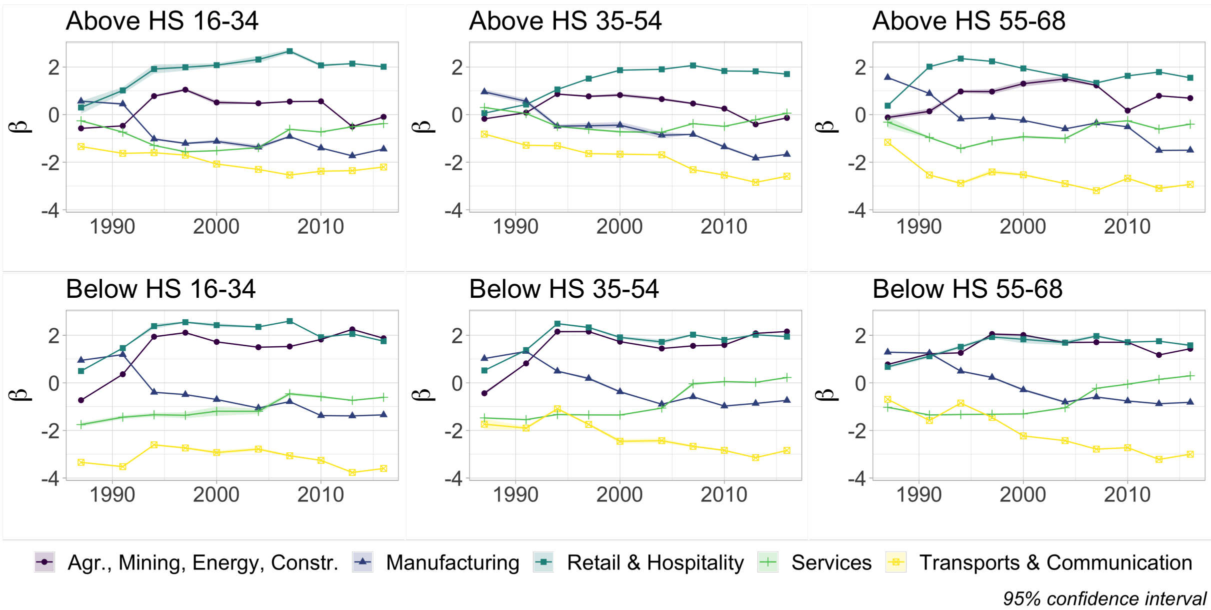

Finally, figure 2.7 presents estimated elasticities of substitution between education level σ, and age groups τ_H and τ_L. τ_L, the elasticity of substitution between non graduates age groups is generally very high, suggesting age groups are perfect substitutes. σ is between 1.68 in 1987-1989 and increases monotonically to 37.67 in 2016-2017. τ_H is between 2.68 in 1987-1989 and 26.50 in 2016-2017.

Figure 2.7: Estimated elasticities of substitution

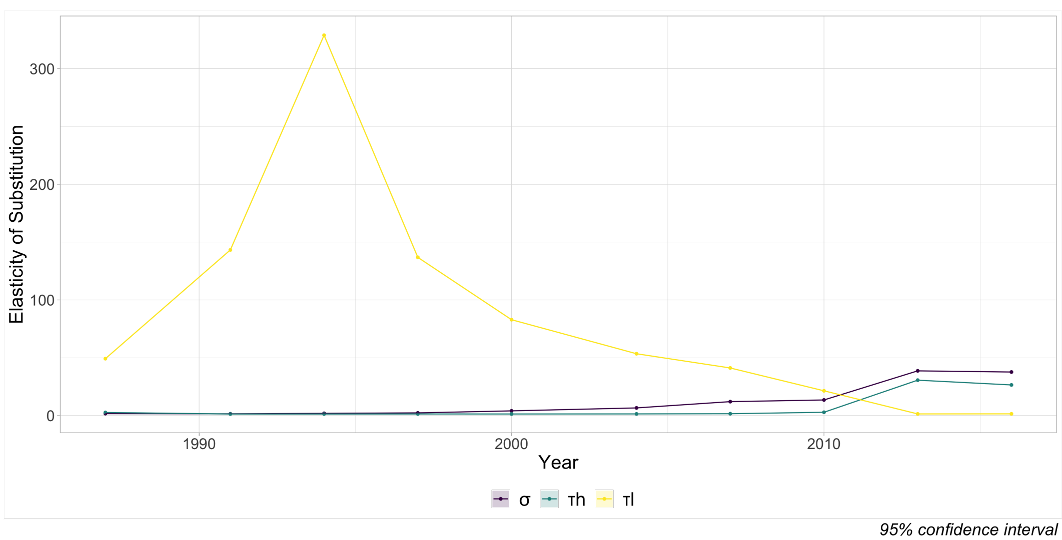

These findings make sense with regards to my estimation method, which unlike most of the literature does not postulate a representative firm by industry, but instead estimates elasticities of substitution at the firm level. Because I account for all firms in the economy and fit the matching distribution, if a large share of firms employs workers of a single education level, either below or above high school, it would be unsurprising to obtain a estimate for σ that suggests that education levels are perfect substitutes (i.e. a very large σ). The same reasoning is valid for substitution between age groups. Scanning the sample for such patterns reveals that a majority of firms employ workers of a single education level: they amount to 78.5% of firms in 1987 and 63.1% in 2017, as well as a single age group: firms who employ a single high school graduates age group amount to 72.8% of firms who employ high school graduates at all in 1987 and 58.2% in 2017. For non graduates, the proportion is between 44.0% in 1987 and 52.0% in 2017. Perfect substitutability of worker types is consistent with the view that the production function at the individual firm level is linear in labor inputs, which is an assumption that had been made in the literature (Hellerstein et al. (1999)).

Discussion. Takeaways from the structural estimates presented in this section are threefold. First, high school graduates productivity has surged over the period. This observation is strongly consistent with a hypothesis of skill-biased technological change, i.e. an increase in worker productivity that favors educated workers, rather than with the competing Hecksher-Ohlin trade hypothesis, which predicts an strengthening of firms' demand for uneducated workers. Second, young and middle-aged high school graduates's share in productivity has increased over the period, which suggests the decreasing high school wage premium for these age groups cannot be explained through an increased demand for experience. Third, workers hold heterogeneous preferences towards sectors. Amenities perceived in the Transport & Communications and Services sector are below zero for most of the period, which puts an upward pressure on wage in these sectors.

These observations must be interpreted in the light of the institutional changes that have occurred in Portugal over the period: the Portuguese labor market is characterized by a steadily (if slowly) increasing minimum wage: in nominal terms, hourly minimum wage is 2.05 euros in 1999 and reaches 3.73 euros in 2017. Most Portuguese workers are also covered by collective bargaining agreements. Finally, it is costly for a firm to fire a worker, because of generous severance packages. Between 2011 and 2014, Portugal has implemented a number of reforms on its labor market: minimum wage was frozen (until the end of 2014), the scope of collective bargaining restricted and terminating workers made less costly. These reforms coincide with a break of trend in the estimated education productivities, and may impact them in two ways. First, the freeze in minimum wage may be partly responsible for the

fact that high school non-graduates productivities stop increasing after 2010: because the minimum wage is binding for a large number of non-graduates, its freeze must reverberate on non-graduates' productivity, since the model predicts it increases wage. However, the minimum wage freeze on its own cannot account for the estimated drop in non-graduates productivity, nor the surge in graduates productivity. Both of these appear to be driven by matching, as the number of firms which employ only high school graduates increases rapidly between 2007 and 2013. In light of the institutional changes that have taken place over these five years, changes in matching have been made easier by lowering workers' severance package and reducing the scope of collective bargaining. Hence an interpretation for the surge in graduates productivity is an overdue increase in firm demand that was kept low before 2010-2013 not because of a low productivity, but because of stringent labor market institutions.

5.2 Model Predictions

Table 2.1 compares the slopes of observed and predicted sorting over time. Slopes are obtained by fitting a time trend to the log of relative education supply in each age group and industry. They can be interpreted as average increase in sorting strength (measured as change in relative supply within an industry) over the period: for instance relative supply increases by on average 127.4% every period in the 16-34 age groups and the primary industries. Model predictions fit the manage to fit the changes in the data quite well, especially for the 16-34 and 35-54 age groups.

Table 2.1: Sorting average yearly percentage growth - Observed versus Predicted

	16-34		35-54		55-68	
Industry	Data	Prediction	Data	Prediction	Data	Prediction
Agr., Mining, Energy, Constr.	1.274	1.037	0.976	0.876	1.178	0.838
Manufacturing	1.478	1.881	1.222	1.774	0.928	1.341
Retail, Hospitality	1.16	1.274	1.315	1.262	0.933	0.909
Transports, Communication	1.391	1.993	1.748	2.248	1.454	1.947
Services	1.03	1.565	1.05	1.273	0.625	0.728
Overall	0.279	0.25	0.226	0.233	0.181	0.253

Figure 2.8 compares observed and predicted average wage by education level and age group over time. Model predictions match the slope of average wages for almost all education levels and age groups, except for senior high school graduates. Average wage for this worker type is over-estimated by the model. This is likely due to the importance of collective

bargaining in Portugal, which presumably tightens the wage distribution and is not accounted for by the model.

Figure 2.8: Average wage by education level and age group - Observed versus Predicted

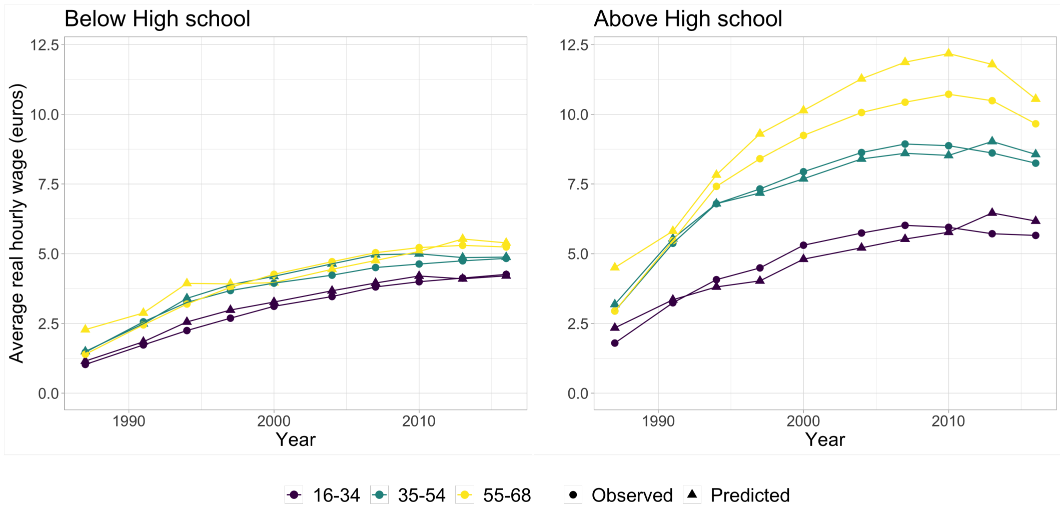

5.3 Counterfactuals

There are four categories of inputs that determine optimal matching and wage and that change over time: the number of workers of each type, the number of firms in each sector, production function parameters and worker preferences parameters. The first two are observed directly in the data and the last two are estimated. In the counterfactuals exercises that follow, I vary each one of the four inputs, holding all other three fixed between 1987-1989 and 2016-2017. The first counterfactual keeps the shares of each sector, production parameters and worker preferences constant to their 1987-1989 levels but lets the worker demography, both in terms of age group and education level, vary as it has in the data between 1987-1989 and 2016-2017. The second counterfactuals holds production parameters, worker preferences and worker demography fixed but lets sector shares vary. The third and fourth counterfactuals vary only production parameters and worker preferences, respectively.

The two object of interests are education-sector sorting and high school wage premium. The model makes predictions on both of these through equilibrium μ and w. Sorting between education and sector is defined as the ratio of high school graduates over non-graduates employed in a sector y, for a given age group a in a given period t:

$$r(\Gamma^t, \beta^t, n^t, m^t) = \frac{\sum_k k_{H,a}\mu_{yk}(\Gamma^t, \beta^t, n^t, m^t)}{\sum_k k_{L,a}\mu_{yk}(\Gamma^t, \beta^t, n^t, m^t)},$$

where μ is the predicted matching. Therefore the change in sorting between two periods t and s is

$$\Delta r_{y,a}^{s,t} = \frac{r_{y,a}(\Gamma^s, \beta^s, n^s, m^s)}{r_{y,a}(\Gamma^t, \beta^t, n^t, m^t)}.$$

Let t=1987-1989 and s=2016-2017. Then the counterfactual change from labor supply is

$$\Delta r_{y,a}^{LB} = \frac{r_{y,a}(\Gamma^t, \beta^t, n^s, m^t)}{r_{y,a}(\Gamma^t, \beta^t, n^t, m^t)}$$

where $r_{y,a}(\Gamma^t, \beta^t, n^s, m^t)$ is the counterfactual sorting if only labor supply n evolves to its 2016-2017 level, while all other factors Γ, β and m stay at their 1987-1989 levels.

Similarly, define wage premium for age group a in a given period t:

$$\omega(\Gamma^t, \beta^t, n^t, m^t) = \frac{\sum_k k_{H,a}\mu_{yk}(\Gamma^t, \beta^t, n^t, m^t) w_{\{H,a\}yk}(\Gamma^t, \beta^t, n^t, m^t)}{\sum_k k_{L,a}\mu_{yk}(\Gamma^t, \beta^t, n^t, m^t) w_{\{L,a\}yk}(\Gamma^t, \beta^t, n^t, m^t)} - 1,$$

so that counterfactual change in wage premium from labor supply is

$$\Delta \omega_{y,a}^{LB} = \frac{\omega_{y,a}(\Gamma^t, \beta^t, n^s, m^t)}{\omega_{y,a}(\Gamma^t, \beta^t, n^t, m^t)}.$$

The next tables 2.2 and 2.3 show the predicted changes in sorting and wage premium along with the four counterfactual scenarios. Note that because of non-linearities, the sum of changes in all four counterfactuals does not sum to the predicted change.

Table 2.2 shows changes in relative employment and the increasing presence of high school graduates in the Manufacturing, Services and Transport & Communications sectors are mainly driven by labor supply: the rise in educated workers' share mechanically increases their employment share in each industry. The counterfactual increase in sorting is uniform across industries however, while predicted sorting is not. An important driver of the heterogeneous increase in industries appear to be the evolution of production parameters and worker preferences: production parameters have a particularly strong positive impact on sorting in Manufacturing and Transport & Communications sectors, while worker preferences drive sorting in Transport & Communications and Services to lower levels than other sectors.

Table 2.2: Changes in Sorting - Predicted versus Counterfactuals

Industry	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
Agr., Mining, Energy, Constr.	3.73	9.4	0.87	0.87	0.77
Manufacturing	22.88	8.01	1.01	1.19	0.76
Retail, Hospitality	8.94	8.51	0.92	0.74	0.85
Transports, Communication	43.33	7.04	0.9	4.12	0.55
Services	14.59	6.4	0.79	0.73	0.61

Interpretation: Predicted relative employment of high school graduates to non-graduates is multiplied by 3.73 between 1987 and 2017 in Agr., Mining, Energy, Constr.

Table 2.3 shows how different scenarios impact changes in wage premium. Consistent with the data, the model predicts a decline in wage premium for all age groups, and especially young workers. Yet each factor except the change in industry composition has a heterogeneous effect on wage premium depending on age group. First labor demographics drive young and middle-aged workers wage premium down. Surprisingly, the same is not true of senior high school graduates: if only labor supply changes over the period, the senior workers' counterfactual wage premium increases. This is because their supply increase, but so does the supply of senior non-graduates, especially relative to younger non-graduates, so that the change in senior worker wage premium is actually positively impacted by labor demographics. The evolution of production parameters has a positive effect on young workers' wage premium, but a negative effect on other age groups, likely because age productivity of middle-aged non-graduates increases and age productivity of senior graduates decreases over the period. Finally workers preferences have a strong, positive impact on all age groups wage premia. Appendix E shows the details of this impact by industry. For young and middle-aged workers, it appears to be mainly driven by the high school wage premium in the Retail & Hospitality and Services sectors.

Table 2.3: Changes in Wage Premium - Predicted versus Counterfactuals

Age group	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
16-34	-0.28	-1.09	-0.7	0.19	3.67
35-54	-0.18	-0.85	-0.72	-0.28	2.23
55-68	-0.01	6.66	-2.14	-1.31	1.07

Interpretation: the predicted wage premium for the 16-34 age group has fallen by 28% between 1987 and 2017

The main takeaway from Table 2.3 is that demographic change, through its impact on labor supply, is the main driver over the decreasing wage premium for young and middle-aged workers. Since estimated parameters show a rise in high school graduates productivity, it seems the prevailing interpretation is that although skill-biased technological change took place over the 1987-2017 period in Portugal, it has been outbalanced by the formidable increase in the relative numbers of high school graduates. This conclusion must be nuanced in the case of senior workers however: their wage premium is mainly dragged down by the changes in industry composition, i.e. structural change, and the evolution of production parameters, as the relative productivity of senior high school graduates with respect to young and middle-aged high school graduates declines, while for non graduates it remains constant.

6 Conclusion

This paper studies wage inequality in Portugal between 1987 and 2017, and seeks to explain the decreasing high school wage premium over the period. The decrease in high school wage premium is particularly stark among young workers, and it is accompanied by a rapidly rising employment share of young educated workers in specific sectors such as Transport & Communications and Services. Over the period, Portugal has experienced a surge in its supply of high school educated workers, as well as sweeping changes in its industry composition, as Services have replaced Manufacturing as the first employer in the country. The increase in educated workers' employment share in the aforementioned sectors suggests a productivity boost in these sectors consistent with skill-biased technological change, but the decreasing high school wage premium observed could also be consistent with a Hecksher-Ohlin theory, whereby less educated workers see their wage rise as they start being more demanded when

Portugal joins the European Union.

To jointly explain changes in sorting between workers and firms, and the decreasing wage premium on the Portuguese labor market, I build a static model of one-to-many matching with transferable utility. Using predictions for both wages and joint distribution of firms and workforces, I am able to separately estimate worker preferences for firms and parameters for firms' nested CES production functions. Estimates show high school graduates productivity has increased in all sectors, consistent with a theory of skill-biased technological change. Counterfactuals suggest changes in sorting are driven by heterogeneity in sectors' relative demand over time, as well as changes in workers' preferences. They also suggest the decreasing high school wage premium is driven mainly by an increase in the relative supply of high school graduates to non-graduates.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.

Almeida, A., Figueiredo, H., Cerejeira, J., Portela, M., Sá, C., and Teixeira, P. N. (2017). Returns to Postgraduate Education in Portugal: Holding on to a Higher Ground? Technical Report 10676, Institute of Labor Economics (IZA). Publication Title: IZA Discussion Papers.

Autor, D., Goldin, C., and Katz, L. F. (2020). Extending the Race between Education and Technology. *AEA Papers and Proceedings*, 110:347–351.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing Inequality: Have Computers Changed the Labor Market? *The Quarterly Journal of Economics*, 113(4):1169–1213.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Azevedo, E. M. and Hatfield, J. W. (2018). Existence of Equilibrium in Large Matching Markets With Complementarities. SSRN Scholarly Paper ID 3268884.

Beaudry, P., Green, D. A., and Sand, B. M. (2015). The Great Reversal in the Demand for Skill and Cognitive Tasks. *Journal of Labor Economics*, 34(S1):S199–S247.

Becker, G. S. (1973). A Theory of Marriage: Part I. *Journal of Political Economy*, 81(4):813–846.

Bikhchandani, S. and Ostroy, J. M. (2002). The Package Assignment Model. *Journal of Economic Theory*, 107(2):377–406.

Blair, P. Q. and Deming, D. J. (2020). Structural Increases in Demand for Skill after the Great Recession. *AEA Papers and Proceedings*, 110:362–365.

Bonhomme, S. (2021). Teams: Heterogeneity, Sorting, and Complementarity. *Working Paper*.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A Distributional Framework for Matched Employer Employee Data. *Econometrica*, 87(3):699–739.

Burdett, K. and Mortensen, D. T. (1998). Wage Differentials, Employer Size, and Unemployment. *International Economic Review*, 39(2):257–273.

Cahuc, P., Postel-Vinay, F., and Robin, J.-M. (2006). Wage Bargaining with On-the-Job Search: Theory and Evidence. *Econometrica*, 74(2):323–364.

Card, D. and DiNardo, J. E. (2002). Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles. *Journal of Labor Economics*, 20(4):733–783.

Card, D. and Lemieux, T. (2001). Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis. *The Quarterly Journal of Economics*, 116(2):705–746.

Cardoso, A. R. (2004). Jobs for Young University Graduates: Is It Worth Having a Degree? Technical Report 1311, Institute of Labor Economics (IZA). Publication Title: IZA Discussion Papers.

Centeno, M. and Novo, A. A. (2014). When supply meets demand: wage inequality in Portugal. *IZA Journal of European Labor Studies*, 3(1):23.

Che, Y.-K., Kim, J., and Kojima, F. (2019). Stable Matching in Large Economies. *Econometrica*, 87(1):65–110.

Chiappori, P.-A., McCann, R. J., and Nesheim, L. P. (2010). Hedonic Price Equilibria, Stable Matching, and Optimal Transport: Equivalence, Topology, and Uniqueness. *Economic Theory*, 42(2):317–354.

Choné, P. and Kramarz, F. (2021). Matching Workers' Skills and Firms' Technologies: From Bundling to Unbundling. *CREST Working Paper*.

Choo, E. and Siow, A. (2006). Who Marries Whom and Why. *Journal of Political Economy*, 114(1):175–201.

Dupuy, A. and Galichon, A. (2014). Personality Traits and the Marriage Market. *Journal of Political Economy*, 122(6):1271–1319.

Dupuy, A. and Galichon, A. (2022). A note on the estimation of job amenities and labor productivity. *Quantitative Economics*, 13(1):153–177.

Ekeland, I., Heckman, J., and Nesheim, L. (2004). Identification and Estimation of Hedonic Models. *Journal of Political Economy*, 112(S1):S60–S109. Publisher: The University of Chicago Press.

Fox, J., Yang, C., and Hsu, D. H. (2018). Unobserved Heterogeneity in Matching Games. *Journal of Political Economy*, Vol 126, No 4.

Fox, J. T. (2010). Identification in matching games. *Quantitative Economics*, 1(2):203–254.

Galichon, A. (2016). *Optimal Transport Methods in Economics*. Princeton University Press.

Galichon, A. and Salanié, B. (2021). Cupid's Invisible Hand: Social Surplus and Identification in Matching Models. *The Review of Economic Studies*.

Goldin, C. and Katz, L. (2008). *The Race Between Education and Technology.* Belknap Press for Harvard University Press.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.

Gretsky, N., Ostroy, J., and Zame, W. (1992). The Nonatomic Assignment Model. *Economic Theory*, 2:103–27.

Heckman, J. J., Matzkin, R. L., and Nesheim, L. (2010). Nonparametric Identification and Estimation of Nonadditive Hedonic Models. *Econometrica*, 78(5):1569–1591.

Hellerstein, J. K., Neumark, D., and Troske, K. R. (1999). Wages, Productivity, and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations. *Journal of Labor Economics*, 17(3):409–446.

Hershbein, B. and Kahn, L. B. (2018). Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings. *American Economic Review*, 108(7):1737–1772.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica*, 87(5):1439–1474.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Kelso, A. S. and Crawford, V. P. (1982). Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica*, 50(6):1483–1504.

Kingma, D. P. and Ba, J. (2017). Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*. arXiv: 1412.6980.

Lindenlaub, I. (2017). Sorting Multidimensional Types: Theory and Application. *The Review of Economic Studies*, 84(2):718–789.

Lise, J. and Postel-Vinay, F. (2020). Multidimensional Skills, Sorting, and Human Capital Accumulation. *American Economic Review*, 110(8):2328–2376.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. *Frontiers in econometrics*.

Sattinger, M. (1979). Differential Rents and the Distribution of Earnings. *Oxford Economic Papers*, 31(1):60–71.

Sattinger, M. (1993). Assignment Models of the Distribution of Earnings. *Journal of Economic Literature*, 31(2):831–880.

Shimer, R. and Smith, L. (2000). Assortative Matching and Search. *Econometrica*, 68(2):343–369.

Valletta, R. G. (2016). Recent Flattening in the Higher Education Wage Premium: Polarization, Skill Downgrading, or Both? NBER Working Paper 22935.

Vohra, R. V. (2011). *Mechanism Design: A Linear Programming Approach*. Econometric Society Monographs. Cambridge University Press, Cambridge.

A Data

I use *Quadros de Pessoal*, a matched employer-employee dataset provided by the Portuguese National Institute (Instituto Nacional de Estatística, INE). *Quadros de Pessoal* is issued yearly from 1987 to 2017, based on firms declarations on their characteristics and their employees'. Both workers and firms are identified across time by a unique identifier.

I use information on firm industry, worker's age and education level Industries are provided as "economic activity", up to 3 digit level. Because of classification changes at the 2 and 3 digits level over time, I use the one digit level classification, to keep consistency over the years. I exclude firms whose economic activity at the 1 digit level are unknown. Worker education is provided as a 3 digits classification, out of which I aggregate 9 levels: no schooling, primary schooling 1 (up to 10 years old), primary schooling 2 (up to 13 years old), primary schooling 3 (up to 15 years old), completed high school, some higher education, bachelor, masters and PhD. Worker age is used directly without further cleaning. I exclude from the sample any worker whose education level of age is unknown (3.9% of observations per year on average)

I also use information on wages and number of hours worked. Wage is provided as a average monthly earnings, that accounts for bonuses and extra hours earnings. Number of hours is provided as the baseline number of hours in the contract, plus any extra hours worked (averaged overt he year). I consider the sum of base and extra hours as my measure for number of hours worked per month. I divide monthly wage by monthly hours to obtain a measure of hourly wage, and deflate it. Real hourly wage is my final measure of wage. I exclude from the sample any worker who has worked zero hours or earned zero wage over the year (11.5% of observations per year on average). These are mainly, in my understanding, workers on sick leave, maternity leave, or sabbatical that do not contribute to firm production in that year. I also exclude from the sample any workers who are strictly under 16 or above 68 (the retirement age in Portugal)

Additionally, I exclude any observation with a missing or 0 worker ID (3.5% of observations per year on average). I am also faced with an issue of duplicate worker IDs which, even though it is minor in the sample later years (about 4.8% of observations per year on average from 2007 to 2017, including 0 IDs), it is much more serious in the earlier years (about 19% of the sample in 1987, including 0 IDs). I suspect these to be encoding mistakes that relate to actual different workers. Some can also be workers who hold two different jobs (for instance

an employee somewhere who also have a self-employed activity). Because I do not use the panel aspect of the data, and therefore encoding mistakes in workers ID are not a problem in my analysis, I keep most duplicates, only removing observations who appear more than 5 times in any given year (an average 6.1% of observations per year, less than 1% of the dataset starting in 2007). I also exclude from the sample any worker who is self-unemployed, in unpaid family care, or labelled under "other" employment contract (7.1% of observations per year on average). The rationale behind not considering self-employed is that many of self-employed workers actually work as consultants for a firm, with no way to link them. Self-employed workers on their own represent about 1% of the dataset.

B Details on Empirical Facts

B.1 Wage levels by gender

Figure 2.9: Average wage by education level and age group - Men

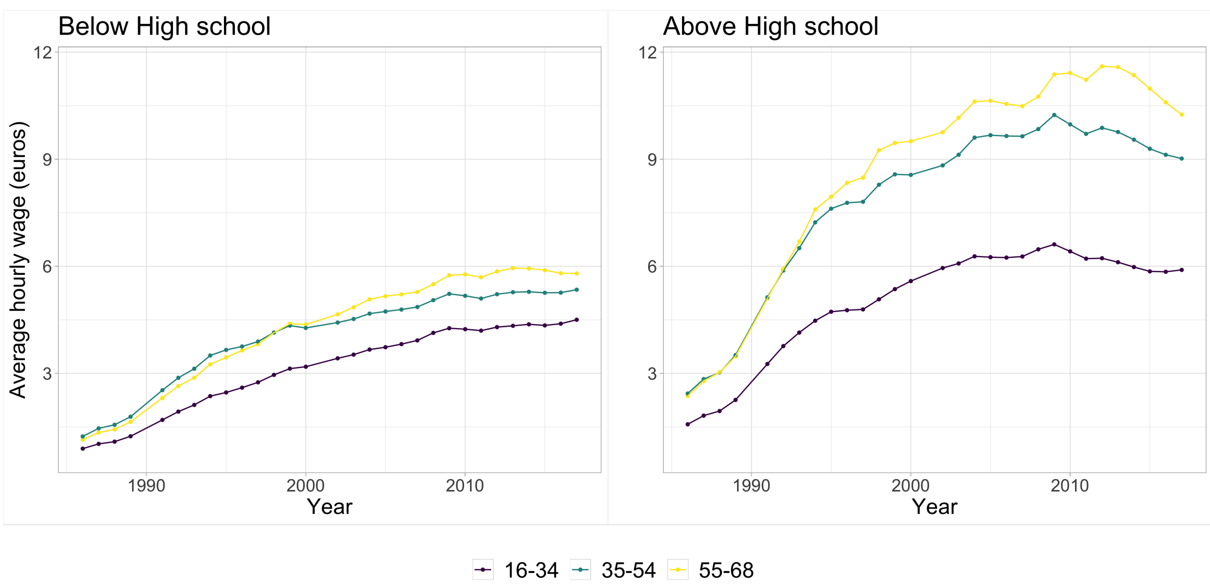

Figure 2.10: Average wage by education level and age group - Women

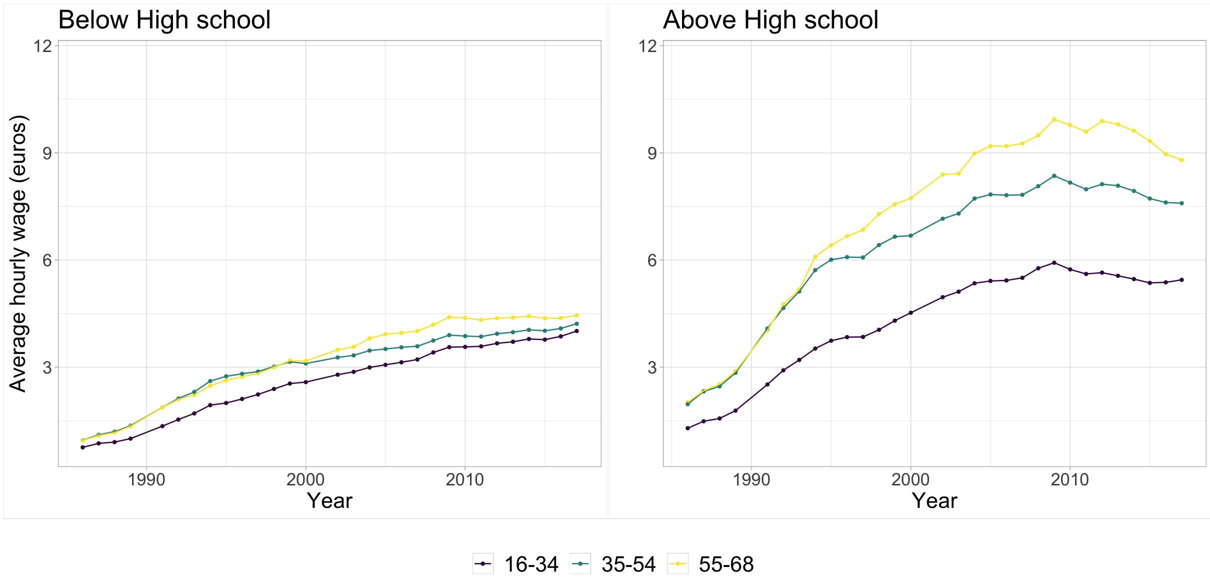

B.2 The university wage premium

The university wage premium is obtained through regressing the following equation:

$$\log w_{ijt} = \sum_{a_i \in \{y,m,s\}} \mathbb{1}_{[\text{University graduate}_i]}\beta_{a_i t} + g_i + r_{jt} + d_{jt} + u_{ijt}.$$

Figure 2.11: University wage premium by age group

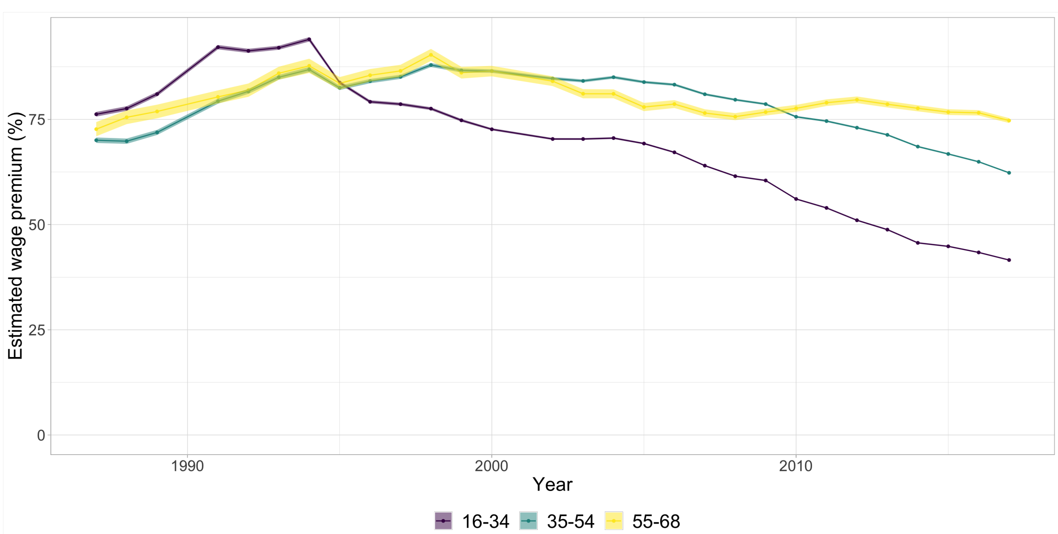

C Proofs

Proposition 1

Proof. Let $Z_1 = \max_{y,k} \{u_{xyk} + \xi\epsilon_{iyk}\}$ and $Z_2 = \max_k \{v_{yk} + \xi\eta_{jk}\}$. The proof consists is showing that Z_1 follows a Gumbel distribution with expectation $\xi \log \sum_y \sum_k \exp\left(\frac{u_{xyk}}{\xi}\right)$ and Z_2 follows a Gumbel distribution with expectation $\xi \log \sum_k \exp\left(\frac{v_{yk}}{\xi}\right)$.

$$\begin{aligned}
\mathbb{P}[Z_1 \le c] =& \mathbb{P}\left[\epsilon_{iyk} \le \frac{c - u_{xyk}}{\xi}\ \forall y, k\right] \\
=& \prod_{y,k} \mathbb{P}\left[\epsilon_{iyk} \le \frac{c - u_{xyk}}{\xi}\right] \\
=& \prod_{y,k} \exp\left(-\exp\left(\frac{u_{xyk} - c}{\xi}\right)\right) \\
\Rightarrow \log \mathbb{P}[Z_1 \le c] =& -\sum_{y,k} \exp\left(\frac{u_{xyk} - c}{\xi}\right) \\
=& -\exp\left(\frac{-c + \log \sum_{y,k} \exp(u_{xyk})}{\xi}\right).
\end{aligned}$$

And a similar reasoning shows:

$$\mathbb{P}[Z_2 \le c] = -\exp\left(\frac{-c + \log \sum_k \exp(v_{yk})}{\xi}\right).$$

Hence up to the Euler-Mascheroni constant, Z_1 follows a Gumbel distribution with expectation $\xi \log \sum_y \sum_k \exp\left(\frac{u_{xyk}}{\xi}\right)$ and Z_2 follows a Gumbel distribution with expectation $\xi \log \sum_k \exp\left(\frac{v_{yk}}{\xi}\right)$. □

Proposition 2

Proof. Following McFadden (1974), Choo and Siow (2006), the probability that worker x chooses option $\bar{y}, \bar{k}$ is

$$\begin{aligned}
\mathbb{P}\left[\bar{y}, \bar{k} = \arg\max u_{xyk} + \xi\epsilon_{yk}\right] &= \mathbb{P}\left[\xi\epsilon_{yk} \le u_{x\bar{y},\bar{k}} - u_{xyk} + \xi\epsilon_{\bar{y}\bar{k}}\ \forall y, k\right] \\
&= \int \prod_{y,k} \exp\left(-\exp\left(\frac{u_{x\bar{y},\bar{k}} - u_{xyk} + \epsilon}{\xi}\right)\right) \exp(-\epsilon) \exp\left(-\exp(-\epsilon)\right)\, d\epsilon \\
&= \frac{\exp\left(\frac{u_{xyk}}{\xi}\right)}{1 + \sum_{y,k} \exp(\frac{u_{xyk}}{\xi})}.
\end{aligned}$$

A similar derivation applied on the firm side. □

Theorem 2.1 Based on Gretsky et al. (1992) and Galichon and Salanié (2021).

Proof. Consider the following problem over the sum of worker welfare $\int_i u_i \, di$ and firm welfare $\int_j v_j \, dj$:

$$\begin{aligned}
&\inf_{u,v} \int_i u_i \, di + \int_j v_j \, dj \\
&\text{s.t} \sum_x \sum_{i:x_i=x}^{k_x} u_i + v_j \geq \Phi_{y_j k} + \xi \sum_x \sum_{i:x_i=x}^{k_x} \epsilon_{iy_j k} + \xi \eta_{jk} \quad \forall k, j \\
&u_i \geq \xi \epsilon_{i0}.
\end{aligned} \tag{2.16}$$

Take any two u, v such that $\sum_x k_x u_{xyk} + v_{yk} \geq \Phi_{yk}$ and $u_{x0} = 0$ and define

$$\begin{cases} u_i = \max_{y,k}\{u_{x_i yk} + \xi \epsilon_{iyk}\} \\ v_j = \max_k\{v_{y_j k} + \xi \eta_{jk}\}. \end{cases}$$

Then (u, v) satisfies (2.16)'s constraints.

Reciprocally, fix any u_i, v_j that satisfy the constraints in this problem and define

Let

$$\begin{cases} u_{xyk} = \min_{i,x_i=x}\{u_i - \xi \epsilon_{iyk}\} \text{ and } u_{x0} = 0 \\ v_{yk} = \min_{j,y_j=y}\{v_j - \xi \eta_{jk}\}. \end{cases}$$

Then the constraint in problem (2.16) becomes $\sum_x k_x u_{xyk} + v_{yk} \geq \Phi_{yk}$.

Applying the law of large numbers, we get that (2.16) is equivalent to

$$\begin{aligned}
&\min_{u,v} \sum_x n_x G_x(u_x) + \sum_y m_y H_y(v_y) \\
&\text{s.t} \sum_x k_x u_{xyk} + v_{yk} = \Phi_{yk} \quad \forall k, y \\
&u_{x0} = 0.
\end{aligned} \tag{2.17}$$

By complementary slackness condition, solving problem (2.16) with $u_{xyk} = \alpha_{xyk} + w_{xyk}$ and $v_{yk} = \gamma_{yk} - \sum_x k_x w_{xyk}$ yields equilibrium wage. Equilibrium supply and demand $S^x_{yk} = k_x D^y_k$ obtain as the Lagrange multiplier μ_{yk} on constraint $\sum_x k_x u_{xyk} + v_{yk} \geq \Phi_{yk}$.

□

Proof. **Theorem 2.2**

Rewrite problem (2.4) as saddle-point:

$$\begin{aligned}
&\min_{u,v}\max_{\mu}\sum_x n_x G_x(u_x)+\sum_y m_y H_y(v_y)\\
&\quad+\sum_{y,k}\mu_{yk}\left(\Phi_{yk}-\sum_x k_x u_{xyk}-v_{yk}\right)+\sum_x S_0^x(-u_{x0})\\
&=\max_{\mu}\sum_{y,k}\mu_{yk}\Phi_{ky}\\
&\quad-\sum_x n_x\max_u\left\{\sum_y \frac{k_x\mu_{yk}}{n_x}u_{xyk}+\frac{S_0^x}{n_x}u_{x0}-G_x(u)\right\}\\
&\quad-\sum_y m_y\max_v\left\{\sum_y \frac{\mu_{yk}}{m_y}v_{yk}-H_y(v)\right\}\\
&=\max_{\mu}\sum_{y,k}\mu_{yk}\Phi_{ky}\,dk\\
&\quad-\xi\left(\sum_x\sum_{y,k}k_x\mu_{yk}\log\frac{k_x\mu_{yk}}{n_x}-\sum_x S_0^x\log\frac{S_0^x}{n_x}-\sum_{y,k}\mu_{yk}\log\frac{\mu_{yk}}{m_y}\right),
\end{aligned}$$

where the last line is obtained through solving for G and H's convex conjugates:

$$\begin{aligned}
G_x^*(\mu)&=\max_u\left\{\sum_{y,k}\frac{k_x\mu_{yk}}{n_x}u_{xyk}+\frac{S_0^x}{n_x}u_{x0}-G_x(u)\right\}\\
H_y^*(\mu)&=\max_v\left\{\sum_{y,k}\frac{\mu_{yk}}{m_y}v_{yk}-H_y(v)\right\}.
\end{aligned}$$

For which first order conditions are

$$\frac{k_x\mu_{yk}}{n_x}=\frac{\exp\left(\frac{u_{xyk}}{\xi}\right)}{\sum_{y,k}\exp\left(\frac{u_{xyk}}{\xi}\right)}\text{ and }\frac{\mu_{yk}}{m_y}=\frac{\exp\left(\frac{v_{yk}}{\xi}\right)}{\sum_k\exp\left(\frac{v_{yk}}{\xi}\right)}.$$

Which ensures that μ is feasible, i.e. satisfies marginal conditions, otherwise the value of the social planner problem is $+\infty$.

Problem (2.5)'s objective function is strictly concave and the maximization set defined by the marginal conditions (2.8) is compact. Therefore the maximum exists and is unique.

$\square$

D Estimation

This section covers the details of log likelihood estimation. Subsection D.1 describes how to data is processed into observed matchings and wages by firm type and workforce, subsection D.2 presents the gradient descent algorithm used for log likelihood maximization and subsection D.3 provides details on log likelihood gradient computation.

D.1 Workforce Discretization

From *Quadros de Pessoal*, I build a workforce matched to each firm every year, by using firm identifiers provided in the data. I weigh workers by their number of monthly hours worked on average over ear, which is directly provided in the dataset. One full-time worker is equivalent to 174 hours worked per month (which is a 40 hours week). If for instance a worker has worked 180 hours per month, she counts as $\frac{180}{174} = 1.03$ full-time workers. The distribution of firms by number of high school graduates and non graduates employed is plotted in figure 2.12 in the periods 1987-1989 and 2016-2017. A firm is defined through a distinct firm identifier-year combination. 2.12 shows a large majority of firms are small firms. Many firms employ no high school graduates, especially at the start of the period: they represent 75.9% of firms in 1987-1989, and 37.0% of firms in 2016-2017. In contrast, firms who do not employ no high school graduates make up 2.6% and 26.4% of all firms, in 1987-1989 and 2016-2017 respectively. Firms who employ more than a thousand of workers at one education level are excluded from the graph, but not from the estimation. They represent 234 firms in 1987-1989 and 206 firms in 2016-2017.

Figure 2.12: Firm distribution by number of high school graduates and non graduates employed

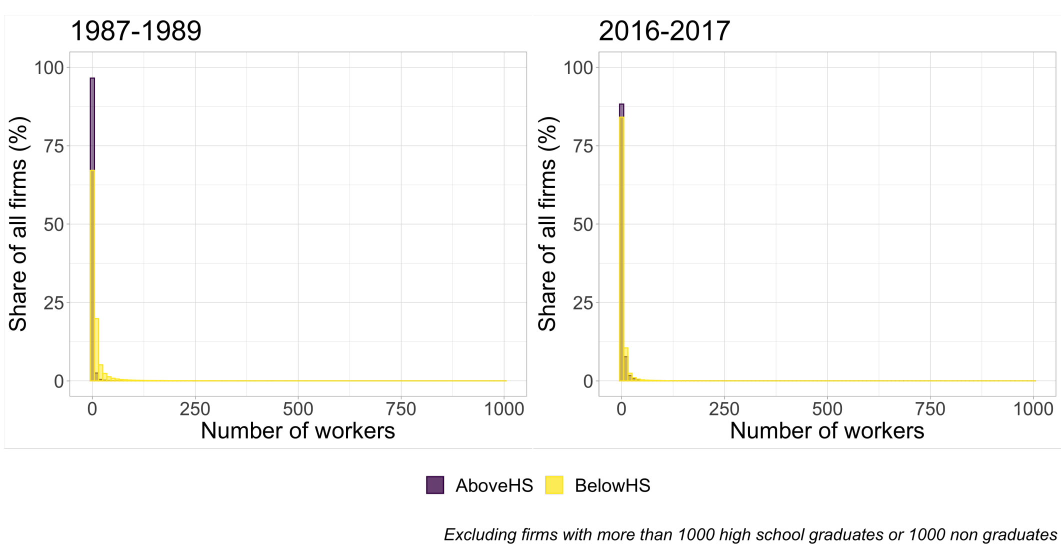

Performing the estimation requires to compute observed matching $\tilde{\mu}_{yk}$ and observed average wage $\tilde{W}_{xyk}$ by workforce k. The number of different observed workforces in the data is very large: there are 96612 combinations in 1987-1989 and 69314 in 2016-2017. Max likelihood computation requires to evaluate μ_{yk} and w_{xyk} on all observed workforces. To speed up the max likelihood computation, I cluster observed workforces into a smaller number of representative workforces. To do so, choose a number of bins B. For each worker type x, split the interval between 0 and k_x^{max} in B smaller intervals, where k_x^{max} is the largest observed number of type x workers employed by a firm. For each worker type x, the procedure yields B intervals, or clusters $[0, k_x^1), \dots, [k_x^{b-1}, k_x^b), \dots, [k_x^B, k_x^{max}]$. Each observed number of worker x employed by a firm falls into one of these intervals. I assign each observed number to a cluster. The representative number of workers for each cluster is $\frac{k_x^b - k_x^{b-1}}{2}$.

In the baseline estimation, $B = 15$. Intervals are split according to a logarithmic scale. The number of observed clusters is reduced to 10359 in 1987-1989 and 21871 in 2016-2017. As an illustration, figure 2.13 displays worker distribution across firms by type, and the clustering of workforce.

Figure 2.13: Worker type distribution and clusters, 1987-1989

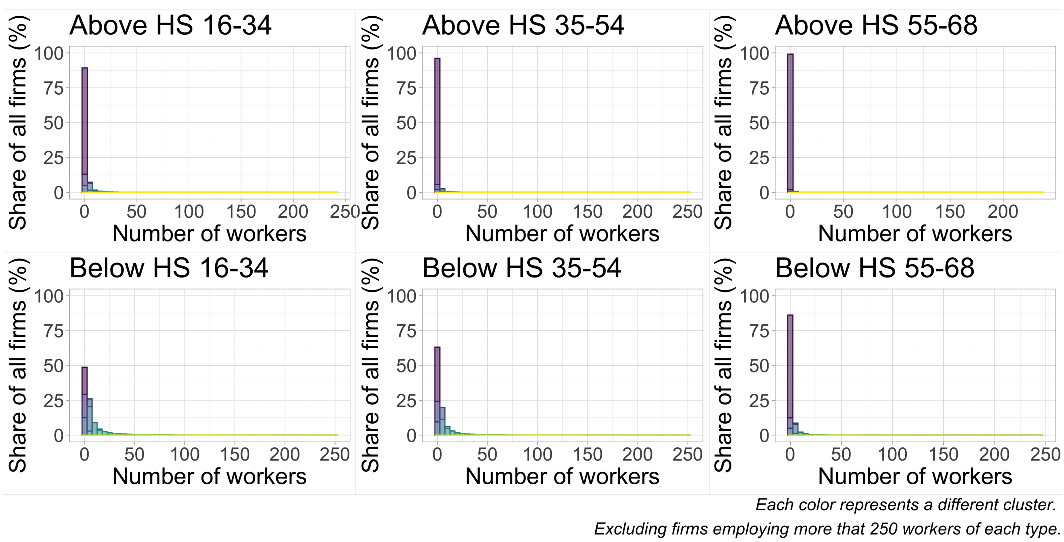

D.2 Adam Algorithm

Adam is a first-order gradient-based optimization algorithm. It belong to the family of algorithms with adaptive learning rates. Their main benefit is speed: they use information given by the gradient to modify their learning rate, and hence improve convergence speed. In particular, Adam uses momentum, i.e. an exponentially moving average of past gradients,

at each iteration. It also uses bias correction. Adam was first introduced by Kingma and Ba (2017). For a general presentation of the algorithm, see Goodfellow et al. (2016). The algorithm applied to the present problem goes as follows:

Set decay rates $\rho_1 = .9$, $\rho_2 = .999$, step $\epsilon = 1e-2$, stabilizer $\delta = 1^{e-8}$ and tolerance $tol = 1^{e-4}$.
Initialize parameters to Γ_0, β_0
Initialize moment variables $s = 0$, $r = 0$ and time step $t = 0$.
While $\max \left| \frac{\nabla_{\Gamma,\beta} l(\Gamma_t, \beta_t, n, m, s_t^2)}{l(\Gamma_t, \beta_t, n, m, s_t^2)} \right| > tol$
Compute $s_t^2 = \frac{1}{W} \sum_x \sum_{y,k} \tilde{K}_{xyk} \left(\tilde{W}_{xyk} - w_{xyk}(\Gamma_t, \beta_t, n, m) \right)^2$
Compute $g \leftarrow \nabla_{\Gamma,\beta} l(\Gamma_t, \beta_t, n, m, s_t^2)$
Update $t \leftarrow t + 1$
Update $s \leftarrow \rho_1 s + (1 - \rho_1) g$ and $r \leftarrow \rho_2 r + (1 - \rho_2) g \odot g$
Correct bias in first moment $\hat{s} \leftarrow \frac{s}{1-\rho_1^t}$ and second moment $\hat{r} \leftarrow \frac{r}{1-\rho_2^t}$
Compute update $\Delta(\Gamma, \beta) = \frac{\hat{s}}{\sqrt{\hat{r}}+\delta}$
Apply update $(\Gamma_{t+1}, \beta_{t+1}) \leftarrow (\Gamma_t, \beta_t) + \epsilon \Delta(\Gamma, \beta)$
end While

D.3 Likelihood gradient

Applying Adam requires to compute likelihood gradient $\nabla_{\Gamma,\beta} l(\Gamma_t, \beta_t, n, m, s_t^2)$. Let $\omega \in (\Gamma, \beta)$ be any of the parameters governing firm production or workers' preferences. Log likelihood differential with respect to ω is

$$\frac{\partial l(\Gamma, \beta, n, m, s^2)}{\partial \omega} = \sum_x \sum_{y,k} k_x \tilde{\mu}_{yk} \frac{\partial \log \mu_{yk}(\Gamma, \beta, n, m, s^2)}{\partial \omega} + \sum_x \tilde{S}_0^x \frac{\partial \log S_0^x(\Gamma, \beta, n, m, s^2)}{\partial \omega} - \sum_x \sum_{y,k} \tilde{K}_{xyk} \frac{\left(\tilde{W}_{xyk} - \frac{\partial w_{xyk}(\Gamma, \beta, n, m, s^2)}{\partial \omega} \right)^2}{2s^2},$$

where

$$\frac{\partial \log \mu_{yk}(\Gamma, \beta, n, m, s^2)}{\partial \omega} = \frac{1}{1 + \sum_x k_x} \left(\frac{\partial \Phi_{yk}}{\partial \omega} - \sum_x k_x \frac{\partial U_x}{\partial \omega} - \frac{\partial V_y}{\partial \omega} \right)$$

$$\frac{\partial \log S_0^x(\Gamma, \beta, n, m, s^2)}{\partial \omega} = -\frac{\partial U_x}{\partial \omega}$$

$$\frac{\partial w_{xyk}(\Gamma, \beta, n, m, s^2)}{\partial \omega} = \frac{1}{1 + \sum_x k_x} \left(\frac{\partial \Phi_{yk}}{\partial \omega} - \sum_x k_x \frac{\partial U_x}{\partial \omega} - \frac{\partial V_y}{\partial \omega} \right) - \frac{\partial \alpha_{xyk}}{\partial \omega} + \frac{\partial U_x}{\partial \omega}$$

$\frac{\partial \Phi_{yk}}{\partial \omega}$ and $\frac{\partial \alpha_{xyk}}{\partial \omega}$ can be computed directly given their assumed functional forms. $\frac{\partial U_x}{\partial \omega}$ and $\frac{\partial V_y}{\partial \omega}$ solve the following linear equations:

$$\sum_{y,k} \frac{k_x}{1+\sum_x k_x} \mu_{yk} \left(\sum_x k_x \frac{\partial U_x}{\partial \omega} + \frac{\partial V_y}{\partial \omega} \right) = \sum_{y,k} \frac{k_x}{1+\sum_x k_x} \mu_{yk} \frac{\partial \Phi_{yk}}{\partial \omega} \quad \forall x$$

$$\sum_{k} \frac{1}{1+\sum_x k_x} \mu_{yk} \left(\sum_x k_x \frac{\partial U_x}{\partial \omega} + \frac{\partial V_y}{\partial \omega} \right) = \sum_{k} \frac{1}{1+\sum_x k_x} \mu_{yk} \frac{\partial \Phi_{yk}}{\partial \omega} \quad \forall y,$$

which are obtained by differentiating marginal conditions (2.8).

E Counterfactuals

Table 2.4: Changes in Sorting in 16-34 age group - Predicted versus Counterfactuals

Industry	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
Agr., Mining, Energy, Constr.	5.18	13	0.81	0.93	0.71
Manufacturing	33.39	12.5	1.03	1.19	0.8
Retail, Hospitality	12.31	10.64	0.9	0.78	0.86
Transports, Communication	44.96	8.09	0.84	1.93	0.46
Services	22.05	9.93	0.74	0.59	0.64

Table 2.5: Changes in Sorting in 35-54 age group - Predicted versus Counterfactuals

Industry	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
Agr., Mining, Energy, Constr.	4.26	11.51	0.91	0.84	0.81
Manufacturing	27.24	8.78	1.01	1.17	0.7
Retail, Hospitality	10.08	10.45	0.93	0.69	0.81
Transports, Communication	65.55	10.82	0.91	5.89	0.52
Services	18.73	7.82	0.81	0.83	0.59

Table 2.6: Changes in Sorting in 55-68 age group - Predicted versus Counterfactuals

Industry	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
Agr., Mining, Energy, Constr.	2.87	7.13	1.33	0.59	0.87
Manufacturing	14.69	2.99	0.86	0.93	0.79
Retail, Hospitality	4.43	6.09	0.76	0.3	0.75
Transports, Communication	32.83	5.53	1.39	28.1	0.53
Services	7.27	2.86	1	0.65	0.58

Table 2.7: Changes in Wage Premium in 16-34 age group - Predicted versus Counterfactuals

Industry	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
Agr., Mining, Energy, Constr.	-0.27	-1.04	-0.64	0.38	-20.76
Manufacturing	-0.56	-1.15	-0.33	0.12	-0.2
Retail, Hospitality	-0.46	-1.03	-0.67	0.14	-4.54
Transports, Communication	0.21	-1.07	-0.52	1.18	0.14
Services	0.71	-1.13	-0.6	0.11	4.47

Table 2.8: Changes in Wage Premium in 35-54 age group - Predicted versus Counterfactuals

Industry	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
Agr., Mining, Energy, Constr.	-0.2	-0.81	-0.6	-0.24	46.31
Manufacturing	-0.15	-1.02	-0.65	-0.27	0.23
Retail, Hospitality	-0.46	-0.78	-0.69	-0.45	39.02
Transports, Communication	0.11	-0.85	-0.61	0.61	0.18
Services	0.81	-0.88	-0.7	-0.07	3.94

Table 2.9: Changes in Wage Premium in 55-68 age group - Predicted versus Counterfactuals

Industry	1987-2017 change	Labor supply	Industry composition	Production parameters	Worker preferences
Agr., Mining, Energy, Constr.	-0.49	5.35	-1.9	-1.13	0.29
Manufacturing	0.6	13.56	-4.65	-1.56	1.71
Retail, Hospitality	-0.17	5.76	-2.35	-1.29	0.99
Transports, Communication	-0.15	5.79	-1.4	-0.89	0.42
Services	0.35	4.82	-1.72	-1.39	1.59

F Comparison to Card & Lemieux's model

Katz and Murphy (1992) and Card and Lemieux (2001) have shown that the CES production function parameters are identified from assuming that labor is optimally supplied to the economy and that wages are competitive, that is assuming that in each year t a representative firm solves

$$\max_{H_a,L_a} \gamma(t) - \sum_{a\in\{y,m,s\}} H_a w_{H,a} - \sum_{a\in\{y,m,s\}} L_a w_{L,a}, \tag{2.18}$$

where $\gamma(t)$ is the CES production function described in section 4 with no dependence on firm type, as in this set up I assume a single representative firm. I also assume in this section that elasticities of substitution τ^H, τ^L, σ, as well as age productivity parameters $(\lambda_{H,a})_a$, $(\lambda_{L,a})_a$ do not vary with time. Wages are competitive and equal to marginal productivity:

$$\begin{aligned} w_{H,a}(t) &= \lambda_{H,a}^{\frac{\tau_H-1}{\tau_H}} H_a(t)^{-\frac{1}{\tau^H}} \times \theta_H(t)^{\frac{\sigma-1}{\sigma}} H(t)^{\frac{1}{\tau_H}-\frac{1}{\sigma}} \times \gamma(t)^{\frac{1}{\sigma}} \quad \forall a \in \{y,m,s\}, \\ w_{L,a}(t) &= \lambda_{L,a}^{\frac{\tau_L-1}{\tau_L}} L_a(t)^{-\frac{1}{\tau^L}} \times \theta_L(t)^{\frac{\sigma-1}{\sigma}} L(t)^{\frac{1}{\tau_L}-\frac{1}{\sigma}} \times \gamma(t)^{\frac{1}{\sigma}} \quad \forall a \in \{y,m,s\}. \end{aligned} \tag{2.19}$$

Which results in relative wage equations:

$$\begin{aligned} \log\left(\frac{w_{H,a}(t)}{w_{H,a'}(t)}\right) &= \frac{\tau_H-1}{\tau_H}\log\left(\frac{\lambda_{H,a}}{\lambda_{H,a'}}\right) - \frac{1}{\tau^H}\log\left(\frac{H_a(t)}{H_{a'}(t)}\right), \\ \log\left(\frac{w_{L,a}(t)}{w_{L,a'}(t)}\right) &= \frac{\tau_L-1}{\tau_L}\log\left(\frac{\lambda_{L,a}}{\lambda_{L,a'}}\right) - \frac{1}{\tau^L}\log\left(\frac{L_a(t)}{L_{a'}(t)}\right). \end{aligned} \tag{2.20}$$

Restricting $(\lambda_{H,a})_a$, $(\lambda_{L,a})_a$'s variation in time, and adding a stochastic shock to account for measurement errors in observed wage and hours worked, relative age productivity and

age elasticities of substitution can therefore be estimated by ordinary least squares through equations:

$$\begin{aligned}\log\left(\frac{w_{H,a}(t)}{w_{H,a_0}(t)}\right) &= d_{H,a,a_0} - \frac{1}{\tau^H}\log\left(\frac{H_a(t)}{H_{a_0}(t)}\right) + u_{H,a,a_0} \\ \log\left(\frac{w_{L,a}(t)}{w_{L,a_0}(t)}\right) &= d_{L,a,a_0} - \frac{1}{\tau^L}\log\left(\frac{L_a(t)}{L_{a_0}(t)}\right) + u_{L,a,a_0},\end{aligned} \tag{2.21}$$

where a_0 is the reference age category. Age productivities $\lambda_{H,a}$, $\lambda_{L,a}$ can then be retrieved from fixed effect d_{H,a,a_0}, d_{L,a,a_0} using normalization conditions (2.11).

Estimates for aggregate labor inputs $H(t)$ and $L(t)$ can be computed from estimated age productivities and elasticities of substitution. First order conditions (2.19) also give an expression for relative wage across education levels:

$$\log\left(\frac{w_{H,a}(t)}{w_{L,a}(t)}\right) - \log\left(\frac{\left(\lambda_{H,a}^{\tau_H-1}\frac{H(t)}{H_a(t)}\right)^{\frac{1}{\tau^H}}}{\left(\lambda_{L,a}^{\tau_L-1}\frac{L(t)}{L_a(t)}\right)^{\frac{1}{\tau^L}}}\right) = \frac{\sigma-1}{\sigma}\log\left(\frac{\theta_H(t)}{\theta_L(t)}\right) - \frac{1}{\sigma}\log\left(\frac{H(t)}{L(t)}\right). \tag{2.22}$$

Assume $\log\left(\frac{\theta_H(t)}{\theta_L(t)}\right)$ follows a linear time trend. Plugging in previously estimated age productivities and elasticities of substitution and adding measurement error gives us equation

$$\log\left(\frac{w_{H,a}(t)}{w_{L,a}(t)}\right) - \hat{f} = l(t) - \frac{1}{\sigma}\log\left(\frac{H(t)}{L(t)}\right) + v_{a,t}, \tag{2.23}$$

where $l(t)$ is a linear function of time and $\hat{f}$ is estimated from equations (2.21).

Weighted Least Square estimation of equations (2.21) and (2.23) are presented in table 2.10 and 2.11. The weights used are the inverse sampling variance of estimated wage gaps[5]. Labor input from any given education level and age bin is computed as the total sum of hours workers per month in a year. Average wage premia between age and within education are used as outcome variable in equation (2.21) and computed yearly and by education level by regressing individual wages on a dummy for age, plus fixed effects for gender, industry and region, to control for composition effects. Average wage premia between education levels and within ages are computed in the same fashion.

[5]In equation (2.23), I weight by the inverse of the sum of the wage gaps and $\hat{f}$ inverse sampling variance

Table 2.10: Estimated age productivities and elasticities of substitution - Reduced Form

	Below High School	Above High School
τ	15.907	15.301
	(2.216)	(2.427)
λ_y	0.332	0.331
	(0)	(0.001)
λ_m	0.333	0.332
	(0)	(0)
λ_s	0.335	0.338
	(0)	(0.001)
R^2	0.994	0.972
Obs.	58	58

Estimated age elasticities of substitution τ in Portugal from 1987 to 2017 are higher than estimates found by Card and Lemieux (2001) for the US, the UK and Canada from the 1970s to the early 1990s, which are between 4 and 6. This reflects the lesser impact of movements in relative age group supply on age group wage differential in Portugal than in the US, UK and Canada. Estimated age productivities are very similar between education levels. They are also balanced between age groups, which suggests no age group is much more productive than another.

Table 2.11: Estimated education productivity growth and elasticity of substitution - Reduced Form

σ	4.933
	(0.151)
$\log \frac{\theta_H}{\theta_L}$	0.018
	(0.001)
R^2	0.974
Obs.	87

Elasticity of substitution between workers below and above high school is also higher in Portugal than what is found by Katz and Murphy (1992) for the US and Card and Lemieux (2001) for the UK and the US, who has estimates between 2 and 2.5. However Card and Lemieux (2001) find no significant effect of relative labor supply on relative wage between education levels in Canada, suggesting a very high substitutability of graduates and non

graduates in that country. their analysis also focuses on college versus high school graduates, which is not directly comparable to my analysis on high school graduates and non graduates, who appear to be more substitutable than college graduates and non graduates. Like Katz and Murphy (1992) and Card and Lemieux (2001), I find evidence of skill-biased technological change in Portugal over the period, as relative productivity between education groups increases by 1.6% every year. This is in the range of what Card and Lemieux (2001) find for the US, UK and Canada.

This analysis informs on the large substitutability of workers between age groups and education levels, as well as the slow but significant high school biased technological change occurring in the Portuguese economy between 1987 and 2017, under simple assumptions on supply and demand. Its conclusion is that it is the increase in relative supply of high school graduates that causes the decrease in wage premium, in particular for young workers, who experience a more important rise in relative supply. 2.14 presents the predicted wage gaps by age group, against observed wage gaps in the data. Coefficients estimated with Card and Lemieux (2001)'s method fail to match well the wage premium evolution over the period: predicted wage is flatter over the period than it is the data, suggesting that imposing linearity in the evolution of relative education productivity and constant age productivity across time restricts the estimation too much.

Figure 2.14: Predicted wage gaps between high school graduates and non-graduates of same age

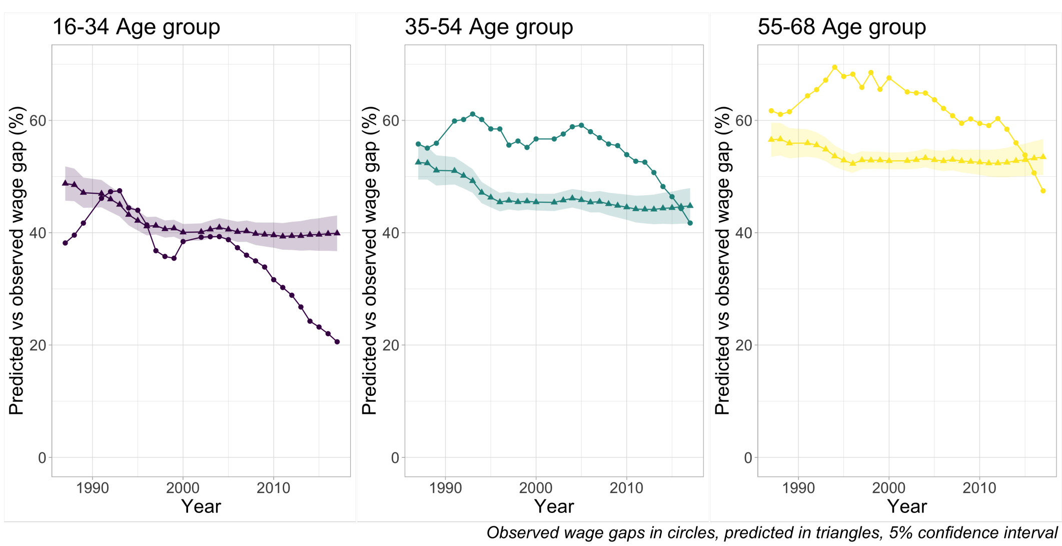

CHAPTER 3

Repeated Matching Games: an Empirical Framework

Chapter co-authored with Jeremy Fox and Alfred Galichon

Abstract

How much does expectations of future returns influence agents matching decisions? To answer this question in the context of labor markets, where workers and firms may both make their employment and hiring decisions depending on future gains, we introduce a model of dynamic matching with transferable utility. Agents have individual states, or types, that evolve depending on current matches. Each period, a matching market with market-clearing prices takes place. We discuss a full equilibrium with time-varying distributions of agents types and show how to computationally solve for it. We introduce econometric shocks to account for unobserved heterogeneity in match formation and show that a stationary equilibrium exists, with and without econometric shocks. We propose two algorithms to compute a stationary equilibrium with econometric shocks, one that solves the stationary equilibrium equations with a non linear solver, the other that reformulates the problem as a min-max problem. We adapt both algorithms for estimation, and use the methods developed to estimate a model of geographic mobility costs for Swedish engineers. We find that mobility costs are impose a sizeable penalty in match production, and evolve non-linearly by age.

1 Introduction

This chapter introduces a tractable model of one-to-one dynamic matching. Agents have individual types, such as education and experience for workers, and industry and occupation for jobs. When deciding with whom to match, agents account for future expected returns that stem from a change in type: for instance, workers expect returns from accumulating different types of human capital depending on which type of jobs (technical, managerial, etc.) they accept. In turn, this change in type will affect returns from future matches. Each period a matching market takes place, where wages act as market-clearing prices.

Relationship formation is an increasingly studied topic in economics, and matching games are a key class of models that predict the formation of relationships. Consequently, matching games have become important empirical tools alongside the availability of datasets on formed relationships. More specifically, economists impose that the observed relationships are the solution to a matching game and then use this assumption to provide restrictions to base estimation on. Combined with an appropriate model of econometric error terms and a computationally tractable estimator, an economist can structurally estimate relevant parameters related to the payoffs agents have for the characteristics of potential partners in the matching game. However, the literature that structurally estimates matching games has almost exclusively, with a few important exceptions mentioned below, restricted attention to static matching games. In a static matching game, agents all make matches (or remain unmatched), and then the model ends. Static matching games capture market forces in the sense that agents compete to make the best matches: for instance, it is hard to get funding or quality venture capital managerial advice for a startup if there are many startups relative to funding opportunities. Likewise, high-quality car part suppliers take profits away from low-quality suppliers. By their nature, static matching games do not model how matches today affect agent state variables and hence future matches. In contrast, many datasets track how agents' relationships change over time. In panel data tracking the personal lives of individuals, one will observe marriage, divorce, and remarriage. In labor market matched employer/employee data, one can track workers as they move between firms and hence track firms as they hire and fire workers. Tech workers in Silicon Valley switching firms may contribute to economic growth. Professional athletes often switch teams, sometimes changing their leagues' competitive landscapes when they do so. In venture capital data, one can see the same startup firm returning for subsequent rounds of funding, perhaps with different sets of venture capitalists each round. Likewise, the same data show venture capital firms investing in different portfolio firms over time. Data on the automotive suppliers providing

car parts to particular car models have a time dimension, as different car models are refreshed, with possibly new suppliers, each year. Key to all these empirical examples is how match partners affect the evolution of an agent's state variable. In labor markets, a worker may gain on-the-job training and hence make better matches in the future. Entrepreneurs might be generalists who require experience in several roles before launching their own firms (Lazear (2009)). In supplier/assembler matching, lower-quality car part suppliers participating in Toyota's Supplier Development Program might raise the quality of future parts (Fox (2018)). For the funding of young startups by venture capitalists, participating in an accelerator program might raise the startup's quality and the prospects of future rounds of funding. In personal relationships, divorce might lead to stigma on the marriage market or it might lead to knowledge of how to avoid relationship mistakes.

To our knowledge, there is not a useful off-the-shelf model from the theory literature that generalizes the static matching games mentioned before to a dynamic setting. Of course, there is a large and influential literature on search models (e.g., Burdett and Mortensen (1998)), but we wish to instead extend static, complete information matching games to a dynamic setting. This proposal introduces a formulation of what we call a repeated matching game. This repeated matching game is a novel model for theoretically understanding the formation of matches over time. Also, the repeated matching game provides a tractable framework for structurally estimating parameters related to agent payoffs using datasets on relationships over time. Our concept of a repeated matching game extends transferable utility, static matching games. In this class of static matching games, agents have complete information about potential partners, and monetary transfers are exchanged between matched agents. The solution concept is often pairwise stability or competitive equilibrium, which can formally coincide for simple matching games. Static, transferable utility matching games have productively formed the basis for many papers that structurally estimate models of relationship formation (Dagsvik (2000), Choo and Siow (2006), Fox (2010b), Chiappori et al. (2017), Fox et al. (2018), Dupuy and Galichon (2014), Galichon and Salanié (2021)).

Our repeated matching game operates in discrete time. Each period, there is a set of active agents. Each agent has a state variable, which is also the type of an agent in the language of static matching games. Making a match, or remaining unmatched can affect the evolution of this agent state variable or agent type. Each period, there is a matching market with prices or transfers for different matches. These prices clear the market. Given these prices, each agent selects the best partner in an, importantly, forward-looking manner. In other words, each agent picks a partner today taking into account how the relationship choice

affects the agent's own state variable and hence the profitability of possibly all matches in future periods. Next period the matching market reopens, new prices are stated and new matches form. Each period should be thought of as long enough for all agents to consider exiting a current match and choosing a new partner. Frictions such as switching costs can be included if desired, for example as one explanation for sticky matches that last multiple periods. A repeated matching game can have both individual and aggregate dynamics. At the individual level, each agent is solving a single-agent dynamic programming problem, where each period the agent's action is to choose a partner to match with. At the aggregate level, the state variable of the matching market is the active agents' current set of types or state variables. This aggregate state variable evolves with the decisions of the individual agents. Matching distribution at the aggregate level is obtained by solving for the social planner's Bellman equation. We expose two methods to compute the social planner's value function and derive equilibrium matching: value function iteration, a cornerstone in solving dynamic optimization problems, and a method derived from the field of deep learning, which uses neural networks to approximate the value function.

One of our most important theoretical results is that a stationary equilibrium exists: there is a set of agent state variables such that, after optimal matches are chosen by forward-looking agents, the same masses of agent types occur. The existence of a stationary equilibrium does not depend on model parameters and lets the researcher optionally ignore aggregate dynamics by imposing that the matching game is at a stationary equilibrium.

A repeated matching game can be a useful empirical framework for structural estimation. We introduce a version of the repeated matching game with econometric errors representing unobserved heterogeneity in the preferences of agents for partner types, following Choo and Siow (2006). We show that a stationary equilibrium also exists in the model with econometric shocks. We then apply two algorithms to compute a stationary equilibrium to the model. One method solves a system of nonlinear equations using a nonlinear programming solver. The second method uses the Chambolle-Pock primal-dual algorithm. We show that both these methods can scale to problems with many agent types. In addition to computing a stationary equilibrium, we can extend the same estimators to structurally estimate parameters in the production of a match with an appropriate dataset. We use a dataset on Swedish engineers moving geographic locations to estimate switching costs by age and geographic distance.

The repeated matching game with econometric errors can best be explained as the combination of two touchstone papers in the literature, although of course many other papers, including ours, are related. Choo and Siow (2006) proposes an estimator for static match-

ing games with logit errors. Rust (1987) proposes an estimator for single agent, dynamic discrete choice models, often using logit errors. In our repeated matching game, an agent's discrete choice each period includes whom to match with and faces logit errors for each type of partner. The agent's type in the matching game is also its state variable, as in dynamic discrete choice models. After computing the prices in a competitive equilibrium, our model of an individual agent's behavior coincides with the dynamic discrete choice model in Rust (1987). In relation with the rest of the literature, We extend classic models of one period, one-to-one, two-sided matching with transferable utility to a setting where such a matching market occurs every time period (Gale (1989), Koopmans and Beckmann (1957), Becker (1973), Shapley and Shubik (1971)). We use econometric assumptions from the literature on estimating static matching games with a continuum of agents (Choo and Siow (2006), Chiappori et al. (2017), Fox et al. (2018), Galichon and Salanié (2021)). Our individual agent problems are dynamic discrete choice models (Miller (1984), Wolpin (1984), Pakes (1986), Rust (1987)). In terms of dynamic matching, Choo (2015) derives closed-form formulas for a model where matched agents are exogenously separated from the pool of agents who can match. By contrast in our models' equilibrium, agents endogenously separate based in part on the availability of attractive partners. Erlinger et al. (2015) and McCann et al. (2015) gave two-period models' equilibrium, where in the first period an agent goes to school and in the second period participates in the labor market. Peski (2021) also focuses on the evolution of individual agent state variables, in his case with a dynamic search model where each period each unmatched agent meets another and accepts or rejects the match. Separations are exogenous and hence unrelated to attractive potential partners, unlike our model.

Section 2 presents the baseline model of repeated matching games and section 3 describes the model with econometric shocks. Section 4 presents our methods for equilibrium computation and section 5 our empirical application. section 6 concludes.

2 The Model Without Econometric Errors

2.1 Agents and Economywide Variables

Let $x \in \mathcal{X}$ be the state of worker, $\mathcal{X}$ is a finite set of worker states. We also call x the type of the worker, recognizing it can change over time. Let $y \in \mathcal{Y}$ be a firm state, with $\mathcal{Y}$ also finite. Both workers and firms have the option to stay unmatched, so that the choice set of workers is $\mathcal{Y}_0$ and the choice set of firms is $\mathcal{X}_0$.

We consider an infinite horizon model, so we drop explicit time subscripts. Note that the horizon is the horizon for the entire economy, rather than the horizon for a worker or firm, which can be finite by placing worker or firm age in the state variables. We discuss below a finite horizon extension for the economy. The worker and firm states evolve according to known transititon rules that are function of the current match (x, y). First $P_{x'|xy} = P(x'|x, y)$ is the conditional probability mass function for the worker state x if matched to a state y firm. Second, $Q_{y'|xy} = Q(y'|x, y)$ is the transition rule for firm state y. For instance, x could track both general work experience (increasing if not unemployed or $y \neq 0$) and occupational-specific experience (increasing for the occupation in the firm state y).

In the aggregate economy, we keep track of the masses of workers and firms of each type. Let m_x be the mass of workers of type x, with $m = (m_x)_{x\in\mathcal{X}}$ being the vector of masses for all worker states. Likewise, let n_y be the mass of firms of type y, with $n = (n_y)_{y\in\mathcal{Y}}$. The aggregate state of the economy is (m, n), which contains the masses of all worker and firm types. Additional macro states, like demand shifters for the industry being studied, can be added to the aggregate state with little conceptual difficulty, although we do not pursue that extension.

In a proposed outcome to the model for one period, let μ_{xy} be the masses of matches between workers of state x and firms of state y. Likewise mu_{x0} is the mass of workers of type x who are unmatched and μ_{0y} is the mass of vacant firms. Let $\mu = (\mu_{xy})_{x\in\mathcal{X}_0, y\in\mathcal{Y}_0}$ be the matrix of matches of masses. In out discussion of estimation, we will have data randomly sampled from μ.

Matched agents exchange monetary transfers in equilibrium. Let w_{xy} be the monetary transfer or wage paid by y to x. let $w = (w_{xy})_{x\in\mathcal{X}, y\in\mathcal{Y}})$ be the matrix of wages for a particular time period. In estimation, we will not use data on monetary transfers, although work on static matching games has explored using data on transfers.

Given that the aggregate state of the economy is (m, n), an outcomes to the model has matches $\mu(m, n)$ and transfers $w(m, n)$ for all possible aggregate states (m, n). The aggregate state transitions using the matches and the individual state transition rules. We use the shorthand notation $(P\mu, Q\mu)$ for next period's aggregate state. We keep the aggregate transition deterministic for simplicity, although adding stochasticity is conceptually straightforward in our framework.

2.2 Individual Agent Problems

There is a common discount factor $\beta < 1$. If a worker of state x matches to a firm of state y at the aggregate state (m, n), the worker receives flow profit

$$\alpha_{xy} + w_{xy}(m, n), \tag{3.1}$$

where α_{xy} is a structural parameter giving the worker flow profit before transfers. If the worker is unmatched, they receive a null wage and we assume $\alpha_{x0} = 0$. The worker is forward looking and chooses a partner y to maximize the expected, present discounted value of lifetime profit, or

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(\alpha_{x^t y^t} + w_{x^t y^t}(m^t, n^t)\right) |x\right], \tag{3.2}$$

where x^t is the worker's state variable in a future period, and y^t is the firm partner type picked that period. Because the individual state transitions are possibly stochastic, future states are random variables. The expectation is over both the future sequence of individual state x and the deterministic evolution of aggregate state (m, n). The worker's problem can be analyzed recursively using his or her Bellman equation:

$$U_x(m, n) = \max_{y \in \mathcal{Y}_0} \left\{ \alpha_{xy} + w_{xy}(m, n) + \beta \sum_{x' \in \mathcal{X}} U_{x'}\left(P\mu, Q\mu\right) P_{x'|xy} \right\}, \tag{3.3}$$

where $U_x(m, n)$ is the continuation value for a worker of individual type x in an economy at aggregate state (m, n). The sum is over next period's individual worker states. Symmetrically, a firm of type y has flow profit

$$\gamma_{xy} - w_{xy}(m, n), \tag{3.4}$$

where γ_{xy} is the non-transfer portion of profit accruing directly to the firm. If the firm is unmatched, it pays no wage and does not produce: $\gamma_{0y} = 0$. The firm's Bellman equation is

$$V_y(m, n) = \max_{x \in \mathcal{X}_0} \left\{ \gamma_{xy} - w_{xy}(m, n) + \beta \sum_{y' \in \mathcal{Y}} V_{y'}\left(P\mu, Q\mu\right) Q_{y'|xy} \right\}, \tag{3.5}$$

where $V_y(m, n)$ is the continuation value of a firm of type y at the aggregate state (m, n).

Each worker and firm is solving a dynamic discrete choice problem, where the discrete choice is a partner type. Other discrete choices, like the decision to undertake an explicit

investment to raise a state variable can be added to the model without changing the basic mathematical structure.

2.3 Competitive Equilibrium

Like the papers on estimating static matching games cited in the introduction, the solution concept for our model is competitive equilibrium. A competitive equilibrium is composed of matches $\mu(m,n)$ and wages $w(m,n)$ for all aggregate states (m,n) such that if $\mu_{xy} > 0$, then the match between worker x and firm y should maximize both agents' profits, as in

$$\mu_{xy} > 0 \Rightarrow \begin{cases} y \in \arg\max_{\tilde{y}\in\mathcal{Y}_0} \left\{\alpha_{x\tilde{y}} + w_{x\tilde{y}}(m,n) + \beta \sum_{x'\in\mathcal{X}} U_{x'}\left(P\mu, Q\mu\right) P_{x'|x\tilde{y}}\right\} \\ x \in \arg\max_{\tilde{x}\in\mathcal{X}_0} \left\{\gamma_{\tilde{x}y} - w_{\tilde{x}y}(m,n) + \beta \sum_{y'\in\mathcal{Y}} V_{y'}\left(P\mu, Q\mu\right) Q_{y'|\tilde{x}y}\right\}, \end{cases} \tag{3.6}$$

where the individual agent's value functions U and V are implicitly computed by value function iteration given the competitive equilibrium. A key result from static matching games like Shapley and Shubik (1971) extends this to repeated matching games.

The decentralized competitive equilibrium satisfies a social planner's problem, due to the transferable utility assumption. The primal problem to maximize the present discounted value of economywide profit given initial aggregate state (m,n), or

$$\max_{\mu^t_{xy}\geq 0} \left\{ \sum_{t=0}^{\infty} \beta^t \sum_{x,y\in\mathcal{X}_0\,\mathcal{Y}_0} \mu^t_{xy}(\alpha_{xy} + \gamma_{xy}) \right\}, \tag{3.7}$$

subject to the feasibility constraints

$$\sum_{y\in\mathcal{Y}_0} \mu^0_{xy} = m_x \quad \forall\, t, x \quad \text{and} \quad \sum_{x\in\mathcal{X}_0} \mu^0_{xy} = n_y \quad \forall\, t, y, \tag{3.8}$$

and the transition rules

$$\sum_{x'y'\in\mathcal{X}\,\mathcal{Y}_0} P_{x|x'y'}\mu^t_{x'y'} = \sum_{y\in\mathcal{Y}_0} \mu^t_{xy} \quad \forall\, t, x \quad \text{and} \quad \sum_{x'y'\in\mathcal{X}_0\,\mathcal{Y}} Q_{y|x'y'}\mu^t_{xy} = \sum_{x\in\mathcal{X}_0} \mu^t_{xy} \quad \forall\, t, y. \tag{3.9}$$

Solving the social planner's primal computes the equilibrium matches $(\mu^t)_t$ for each starting aggregate state (m,n). One can also derive the dual problem, which allows calculation of the equilibrium monetary transfers. The Bellman equation for the social planner's dual

problem at aggregate state (m, n) is

$$W(m,n) = \min_{U_x, V_y} \left\{ \sum_{x \in \mathcal{X}} m_x U_x^0 + \sum_{y \in \mathcal{Y}} n_y V_y^0 \right\}, \tag{3.10}$$

subject to the pairwise stability constraints

$$\begin{aligned} U_x^t + V_y^t \geq & \ (\alpha_{xy} + \gamma_{xy}) \\ & + \beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1}(P\mu, Q\mu) P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1}(P\mu, Q\mu) Q_{y'|xy} \quad \forall\, x, y, t \\ U_x^t \geq & \beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1}(P\mu, Q\mu) P_{x'|x0} \quad \forall\, x, t \\ V_y^t \geq & \beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1}(P\mu, Q\mu) Q_{y'|0y} \quad \forall\, y, t. \end{aligned} \tag{3.11}$$

Once the equilibrium present discounted value of lifetime profit U_x and V_y are computed for all worker and firm types, a set of equilibrium transfers $w(m, n)$ for all (m, n) can be computed.

The result of equivalence between the decentralized equilibrium and the social planner's primal problem is given by the following.

Theorem 2.1. *The matching policy $\mu(m, n)$ in a competitive equilibrium maximizes a social planner's primal problem.*

Proof. In appendix A. □

The primal problem can be analyzed recursively with the social planner's Bellman equation

$$W(m,n) = \max_{\mu_{xy} \geq 0} \left\{ \sum_{xy \in \mathcal{X}_0 \mathcal{Y}_0} \mu_{xy} \left(\alpha_{xy} + \gamma_{xy}\right) + \beta W\left(P\mu, Q\mu\right) \right\} \tag{3.12}$$

subject to the constraint

$$\sum_{y \in \mathcal{Y}_0} \mu_{xy} = m_x \quad \forall\, x \quad \text{and} \quad \sum_{x \in \mathcal{X}_0} \mu_{xy} = n_y \quad \forall\, y. \tag{3.13}$$

The right side of Bellman equation (3.12) is a contraction, so a unique present discounted value for economy wide profit on each aggregate state (m, n) exists across all equilibria.

Theorem 2.2. *A competitive equilibrium exists and the economywide sum of future profits $W(m, n)$ is uniquely determined.*

Proof. In appendix A. □

Note that in many parametrizations, the matches $\mu(m, n)$ are also uniquely determined across all competitive equilibria.

Typically, the aggregate state (m, n) varies from period to period. The time series (m, n) is deterministic, given a starting value for the aggregate state. The entire competitive equilibrium depends on the model parameters: α, γ, β, P and Q. Inspecting the above notation indicates that only the sums $\alpha_{xy} + \gamma_{xy}$ matter, so we can restrict attention to the flow production of a match, defined to be

$$\Phi_{xy} = \alpha_{xy} + \gamma_{xy}. \tag{3.14}$$

2.4 Stationary Equilibrium

Define a constant aggregate state as an aggregate state (m, n) such that in a competitive equilibrium

$$m = P\mu(m, n) \quad \text{and} \quad n = Q\mu(m, n). \tag{3.15}$$

In other words, a constant aggregate state (m, n) remains at the value (m, n) forever after that state is reached. A stationary equilibrium is then defined to be a constant aggregate state (m, n) and the corresponding competitive equilibrium (μ, w). The theoretical result for existence follows.

Theorem 2.3. *For any total mass of agents M, a stationary equilibrium exists.*

Proof. In appendix A. □

Corollary 1. *For any total mass of agents M, an aggregate state exists.*

Working with time-varying aggregate state or restricting attention to a stationary equilibrium is a modeling decision of the researcher. By assuming that the model is at stationary equilibrium, firms will not be concerned about how the distribution of worker types changes over time. The entire focus will be on the evolution of state variables such as the experience levels of individual workers and firms. Other researchers may wish to study the aggregate dynamics of the repeated matching game. Workers and firms may indeed track the overall distribution of agent types. In macro, models with both heterogeneous agents and aggregate dynamics are common (Rios-Rull (1995), Krusell and Smith (1998)). In the dynamic games used in industrial organization, agents best respond to all other agents that track their own states and best respond to an analog to a constant aggregate state (Weintraub

et al. (2008)). As stated before, one can include additional state variables, like industrywide demand shifters, in the repeated matching game framework.

Real datasets often have agents entering and leaving a matching market. For both labor and marriage markets, agents become adults and later either retire or die, exiting the matching market. For venture capital markets, new startups enter and may permanently exit each year. Modeling an exogenous process for entry and exit adds some notational complexity but is conceptually easy to add to the previous description of the repeated matching game. Note that while in some cases (say a worker needing to retire) the horizon of an agent is finite, the overall repeated matching game still has an infinite horizon.

2.5 Finite Horizon for the Economy

It is straightforward to have individual workers and firms with finite horizons. Adding a finite horizon for the economy changes how one solves the primal problem. The social planner problem in (3.7) becomes, with a finite number of periods T

$$\max_{\mu^t_{xy} \geq 0} \left\{ \sum_{t=0}^{T} \beta^t \sum_{x,y \in \mathcal{X}_0 \mathcal{Y}_0} \mu^t_{xy} (\alpha_{xy} + \gamma_{xy}) \right\}. \tag{3.16}$$

Using computational experiments, we have found that it is best to solve this finite-horizon planners' problem directly as written, rather than writing the planner's problem recursively using Bellman's equation. There is not likely to be a stationary equilibrium for a finite horizon problem.

3 The Model With Econometric Errors

The previous model will often predict that some matches occur with mass of zero, meaning $\mu_{xy} = 0$ for some types x and y. This contradicts available datasets where, with enough observations, it may be the case that μ_{xy} is rarely or never zero. This contradiction is solved by accounting for an econometric error terms in the flow profits of both workers and firms. This error term covers variables that matter for match formation, but are unobserved to the econometrician. In this section, we focus on a stationary equilibrium, for simplicity, although the full model with time-varying aggregate states can also be extended to have econometric errors.

3.1 Econometric Preference Shock

Let the flow profit for a worker i of type x be

$$\alpha_{xy} + w_{xy} + \epsilon_{iy}, \tag{3.17}$$

where ϵ_{iy} is worker i's preference shock for type y partners. Worker i is indifferent between all partners of the same observed type y. Likewise, let the flow profit for a firm j of type y be

$$\gamma_{xy} - w_{xy} + \eta_{xj}, \tag{3.18}$$

where η_{xj} is firm j's preference for workers of type x. We make similar assumptions to Choo and Siow (2006) for static matching games and Rust (1987) for single-agent dynamic discrete choice models.

Assumption 3.1. *The econometric errors satisfy the following assumptions:*

1. *For each pair of two workers i and i', $(\epsilon_{iy})_{y\in\mathcal{Y}_0}$ and $(\epsilon_{i'y})_{y\in\mathcal{Y}_0}$ are mutually independent in every time period and across time period. The same mutual independence condition holds for a pair of firms or a pair of one worker and one firm.*
2. *For a single worker i in the two time periods t and $t+1$ with measured states x_i^t and x_i^{t+1}, the distribution of $\left(\epsilon_{iy}^{t+1}\right)_{y\in\mathcal{Y}_0}$ satisfies the following conditional independence, meaning*
$$F\left(\left(\epsilon_{iy}^{t+1}\right)_{y\in\mathcal{Y}_0}\middle| x_i^t, x_i^{t+1}, \left(\epsilon_{iy}^{t}\right)_{y\in\mathcal{Y}_0}\right) = F\left(\left(\epsilon_{iy}^{t+1}\right)_{y\in\mathcal{Y}_0}\middle| x_i^{t+1}\right). \tag{3.19}$$
A similar conditional independence assumption holds for firms.

In other words, knowing one's own preferences for measured partners types is not information about other agents' preferences. Under this assumption, it is irrelevant whether the preference shocks ϵ_{iy} and η_{xj} are public or private information to other players. Also, the assumption states that preferences are drawn anew each time period conditional on measured states x or y, rather than being statistically dependent over time. Relaxing conditional independence from Rust (1987) can be done in several ways, including allowing for time-invariant unobserved types or using instrumental variables, as in Berry and Compiani (2020). We do not formally explore weakening Assumption 3.1 in this paper. Unmeasured preferences in the literature on estimating static matching games with a small number of matching markets, each with a continuum of age,ts, are typically preferences over measured partner types x or y rather than unmeasured preferences attributes (Choo and Siow (2006), Dupuy and Galichon (2014), Chiappori et al. (2017), Fox (2018), Galichon and Salanié (2021)). This

contrasts with a data scheme of many smaller markets, where agents could have preferences over unmeasured (in data) attributes of partners (Fox et al. (2018)).

3.2 Competitive Equilibrium With Econometric Errors

Like in the model without econometric errors (or shocks/heterogeneity), we show that the matching policy in the competitive equilibrium with shocks maximizes a social planner problem which writes as

$$\max_{\mu^t_{xy}\geq 0}\left\{\sum_{t=0}^{\infty}\beta^t\left(\sum_{x,y\in\mathcal{X}_0\,\mathcal{Y}_0}\mu^t_{xy}\Phi_{xy}-\mathcal{E}(\mu^t,n^t,m^t)\right)\text{ s.t (3.9)}\right\}. \tag{3.20}$$

The term $\mathcal{E}$ is the expectation of the econometric errors. It is referred to as the entropy (a term borrowed from the optimal transport literature). Its final form depends on the distributional assumptions on ϵ and η. The social planner problem is subject to the same feasibility constraint and transition rule as in the model without heterogeneity. However, feasibility constraints (3.8) need not appear explicitly: the entropy ensures they are satisfied. if not, the value of entropy goes to infinity. The problem also rewrites as a social planner Bellman equation:

$$W(m,n)=\max_{\mu_{xy}\geq 0}\left\{\sum_{xy\in\mathcal{X}_0\,\mathcal{Y}_0}\mu_{xy}\Phi_{xy}-\mathcal{E}(\mu,n,m)+\beta W\left(P\mu,Q\mu\right)\right\}. \tag{3.21}$$

Define expected indirect payoffs G_x and H_y

$$G_x(u^t)=\mathbb{E}\left[\max_{y\in\mathcal{Y}_0}\left\{u^t_{xy}+\epsilon_y\right\}\right]\quad\text{and}\quad H_y(v^t)=\mathbb{E}\left[\max_{x\in\mathcal{X}_0}\left\{v^t_{xy}+\eta_x\right\}\right] \tag{3.22}$$

and total expected indirect payoffs

$$G(u^t,m^t)=\sum_{x\in\mathcal{X}}m^t_xG_x(u^t)\quad\text{and}\quad H(v^t,n^t)=\sum_{y\in\mathcal{Y}}n^t_yH_y(v^t). \tag{3.23}$$

Then the entropy in social planner problem (3.20) is

$$\mathcal{E}(\mu^t,m^t,n^t)=G^*(\mu^t,m^t)+H^*(\mu^t,n^t), \tag{3.24}$$

where $(G^t)^*$ and $(H^t)^*$ are the entropy of choices:

$$\begin{aligned} G^*(\mu^t, m^t) &= \max_u \left\{ \sum_{xy \in \mathcal{X}_0 \mathcal{Y}} \mu^t_{xy} u_{xy} - G(u, m^t) \right\} \\ H^*(\mu^t, n^t) &= \max_v \left\{ \sum_{xy \in \mathcal{X} \mathcal{Y}_0} \mu^t_{xy} v_{xy} - H(v, n^t) \right\}. \end{aligned} \tag{3.25}$$

The dual to social planner problem (3.20) writes

$$\min_{u^t_{xy}, v^t_{xy}} \sum_{x \in \mathcal{X}} m^1_x G_x(u^1) + \sum_{y \in \mathcal{Y}} n^1_y H_y(v^1) \tag{3.26}$$

subject to the pairwise stability constraints

$$\begin{aligned} u^t_{xy} + v^t_{xy} &\geq \Phi_{xy} + \beta \sum_{x' \in \mathcal{X}} G_{x'} \left(u^{t+1}\right) P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} H_{y'} \left(v^{t+1}\right) Q_{y'|xy} \quad \forall\, t, x, y \\ u^t_{x0} &\geq \beta \sum_{x' \in \mathcal{X}} G_{x'} \left(u^{t+1}\right) P_{x'|x0} \quad \forall\, x \\ v^t_{0y} &\geq \beta \sum_{y' \in \mathcal{Y}} H_{y'} \left(v^{t+1}\right) Q_{y'|0y} \quad \forall\, y \end{aligned} \tag{3.27}$$

The model with heterogeneity affords the same result as without heterogeneity: the competitive equilibrium can be found through solving the social planner problem.

Theorem 3.1. *The matching policy in a competitive equilibrium with heterogeneity maximizes social planner problem* (3.20) *subject to* (3.8) *and* (3.9)*.*

Proof. In appendix A. □

Writing the model with heterogeneity lets us express the competitive matching policy in closed form, provided we assume a distribution for the econometric errors. We follow the literature, and in particular Choo and Siow (2006) and Rust (1987) and assume the following

Assumption 3.2. *Econometric errors ϵ and η are Extreme value 1 distributed.*

Under Assumption 3.2, expected indirect payoffs have logit form (Galichon and Salanié (2021)):

$$G_x(u^t) = \log \sum_{y \in \mathcal{Y}_0} \exp(u^t_{xy}) \quad \text{and} \quad H_y(v^t) = \log \sum_{x \in \mathcal{X}_0} \exp(v^t_{xy}) \tag{3.28}$$

and the entropy $\mathcal{E}$ is

$$\mathcal{E}(\mu, m, n) = \sum_{xy \in \mathcal{X} \mathcal{Y}_0} \mu_{xy} \log \frac{\mu_{xy}}{m_x} + \sum_{xy \in \mathcal{X}_0 \mathcal{Y}} \mu_{xy} \log \frac{\mu_{xy}}{n_y}. \tag{3.29}$$

Proposition 1. *Under Assumption 3.2 competitive matching* μ^t *is*

$$\mu^t_{xy} = \sqrt{m_x n_y} \exp\left(\frac{\Phi_{xy} + \beta \sum_{x' \in \mathcal{X}} U^{t+1}_x P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} V^{t+1}_y Q_{y'|xy} - U^t_x - V^t_y}{2}\right)$$
$$\mu^t_{x0} = m_x \exp\left(\beta \sum_{x' \in \mathcal{X}} U^{t+1}_x P_{x'|xy} - U^t_x\right) \tag{3.30}$$
$$\mu^t_{0y} = n_y \exp\left(\beta \sum_{y' \in \mathcal{Y}} V^{t+1}_y Q_{y'|xy} - V^t_y\right),$$

where (U^t, V^t) *and* (U^{t+1}, V^{t+1}) *are the expected indirect payoffs in period* t *and* $t+1$*.*

Proof. $(\mu)_{xy \in \mathcal{X}_0 \mathcal{Y}_0}$ derives from first order conditions of social planner problem (3.20). □

3.3 Stationary Equilibrium With Gumbel Errors

Assuming econometric errors are Extreme Value 1, as in (3.2), we are able to show that a stationary equilibrium exists for any total mass of agents M. The result is the same as in the model without heterogeneity, although the method of proof is different, as is detailled in appendix A. We do not explore the existence of stationary equilibria outside the logit case in this paper, although it may be that a more general proof exists for any distribution of ϵ and η.

In the model with Gumbel econometric shocks, aggregate state (m, n) is constant and associated matching policy μ is part of a stationary equilibrium if and only if

$$\begin{gathered} m = P\mu \quad \text{and} \quad n = Q\mu \\ (\mu, U, V, U', V') \text{ satisfy relation (3.30)} \\ U = U' \text{ and } V = V' \end{gathered} \tag{3.31}$$

Note that in a stationary equilibrium, indirect expected payoffs to each agent type are the same in every period, and equal to the Lagrange multiplier of the stationarity constraints, up to a constant.

Using (**??**), we are able to establish the following existence result.

Theorem 3.2. *For any total mass of agents M, a stationary equilibrium and aggregate state exist in the model with Gumbel econometric errors.*

Proof. In appendix A □

As in the model without heterogeneity, a stationary equilibrium and aggregate state exist for a given total mass of agents. In the next section, we present different methods to compute both the equilibrium with aggregate dynamics, and the stationary equilibrium.

4 Methods for Equilibrium Computation

This section develops methods for computing equilibria both in a non-stationary and stationary environment, given transition rules and a total mass of agents. In the non-stationary environment, we rely on value function iteration, making use of the social planner Bellman equation (3.21). We also develop a method that builds on neural networks. In the stationary environment, we develop two methods to compute the equilibrium: one uses the Mathematical Programming with Equilibrium Constraint (MPEC, see Su and Judd (2012)) formulation of our problem, and the other rewrites the stationary equilibrium equations as a min-max problem and solves it using techniques from convex optimization (Chambolle and Pock (2011)). We also show how to adapt these methods to estimation using data on matches, assuming that the data represent a stationary equilibrium.

4.1 Equilibrium with Aggregate Dynamics

Value Function Iteration

Equilibrium in dynamic models is classically found using value function iteration (VFI) on the social planner's Bellman equation. We expose it here in our model with econometric errors, but it also applies in the case without heterogeneity. Value function iteration operates on a grid of aggregate states $((m_g, n_g))_{g\in\{1,\dots,G\}}$, where G is the chosen number of points in the grid, and chooses an initial W^0 on each point of this grid. It then proceeds to update W^t as follows

$$\begin{aligned} & W^{t+1}(m_g, n_g) = T\left(W(m_g, n_g)\right) \quad \forall g, \\ \text{where} \quad & T\left(W(m,n)\right) = \max_{\mu_{xy}\geq 0}\left\{\sum_{xy\in\mathcal{X}_0\mathcal{Y}_0} \mu_{xy}\Phi_{xy} - \mathcal{E}(\mu, n, m) + \beta W\left(P\mu, Q\mu\right)\right\}/ \end{aligned} \tag{3.32}$$

Because map T is a contraction, $(W^t)_t$ eventually converges to the fixed point of the social planner's Bellman equation (3.21).

Note that because $(P\mu, Q\mu)$ does not necessarily land on a point of grid $((m_g, n_g))_{g\in\{1,\dots,G\}}$, some interpolation technique is needed to compute the value of W^t at this point: we can use

polynomial interpolation with either simple or Chebyshev polynomials.

Value function iteration is a robust way of computing the equilibrium with aggregate dynamics: once a fixed point W has been found, the competitive equilibrium can be computed on any point of the grid. Its main drawback is that it suffers from a curse of dimensionality: the larger G, the longer each iteration, because to update from W^t to W^{t+1} the optimal matching policy μ must be computed on every point of the grid. Ideally G would be as large as possible to obtain a precise value of W. Within this 'outer' curse of dimensionality, lies an 'inner' curse: the larger X and Y, the number of types of workers and firms, the longer each individual maximization on the grid will take, as the dimension of μ increases.

These two issues can be alleviated in two ways: the 'inner' issue can be tackled using recent optimization solvers to solve the maximization over μ. We use the Nlopt package, which is available in a wide array of programming languages, and find that the maximization on each point of the grid is solved quickly. Besides, because each maximization problem is independent, the loop running through all grid points to solve them can be parallelized. Parallelization partly solves the 'outer' curse of dimensionality. Finally, the numerical analysis literature (Fang and Saad (2009), Walker and Ni (2011)) provides a wide array of methods to accelerate fixed-point iterations. In our own value function iteration, we use the Anderson acceleration method. Its main idea is to use not only W^t to update to W^{t+1}, but also W^{t-1}, W^{t-2}, ... up to some threshold m decided by the analyst. Formally, Anderson acceleration writes

$$W^{t+1} = W^t - \nabla^t f^t \tag{3.33}$$

where[1]

$$\begin{cases} f^t = T(W^t) - W^t \\ \nabla^t = -I + (\mathcal{W}^t + \mathcal{F}^t)\left(\left(\mathcal{F}^t\right)^\top \mathcal{F}^t\right)^{-1}\left(\mathcal{F}^t\right)^\top \\ \mathcal{W}^t = (W^{t-m+1} - W^{t-m}, \ldots, W^t - W^{t-1}) \\ \mathcal{F}^t = (f^{t-m+1} - f^{t-m}, \ldots, f^t - f^{t-1}). \end{cases} \tag{3.34}$$

Note the similarity of Anderson's acceleration with a quasi Newton descent, in which the hessian matrix is approximated by the inverse of the Jacobian of $f(W) = T(W) - W$. See Walker and Ni (2011) for an extended discussion of this point.

[1] As long as $t < m$, the definition for $\mathcal{W}^t$ and $\mathcal{F}^t$ is adapted to include all terms since the first iteration.

Artificial Neural Networks

Instead of using value function iteration, which relies in a grid and is therefore subject to a curse of dimensionality, one can also use the deep learning and neural networks literature[2] to find a fixed point to T, where T is defined as in (3.32). This idea somewhat alleviates the curse of dimensionality, because it does not rely on a grid. It also performs better in terms of accuracy on off-grid points. It is starting to be used in economics with promising results (see for instance Barany and Holzheu (2021)).

Let us draw a sample of masses of size N: $((m_i, n_i))_{i \in \{1,\dots,N\}}$. Let $ANN(m_i, n_i|\theta)$ be the neural network representation of $W((m_i, n_i)$, given vector of parameters θ. Then to find a fixed point of T, one can minimize the following loss function on the sample

$$\min_{\theta} \frac{1}{N} \sum_{i} \left(ANN(m_i, n_i|\theta) - T(ANN(m_i, n_i|\theta)) \right)^2 . \tag{3.35}$$

Such a minimization problem is readily implemented in many libraries in Python (Pytorch, TensorFlow) and Julia (Flux).

Comparison

Using Anderson acceleration on the value function iteration brings significant speed gains: in a model with two types on each side of the market, 'classic' value function iteration takes more than 5 hours on a 6 Cores laptop, while Anderson value function iteration takes only 93 minutes. In both methods, simple polynomial are used for interpolation between grid points and maximizations on grid points are parallelized. Comparison between value function iteration and deep learning is less straightforward because the methods differ on several levels. First the 'procedure' is different: VFI iterates on the grid until it has found a fixed point to the value function, while the deep learning method aims at finding the fixed point by minimizing a loss function. Second, the rely on different types of grid: VFI rests on a rectangular grid and a given family of polynomials for interpolation (simple, Chebyshev, etc.), while deep learning uses a sample of points and neural networks for interpolation.The choice of polynomial family for VFI and neural network for deep learning impacts both the speed of convergence, and the out-of-sample precision. Finally, the two methods are implemented using different tools: value function iteration can be accelerated using parallelization on the computer's cores, while estimation of artificial neural networks can be sped up using a

[2]This section owes to Julien Pascal's excellent blog post on the topic of artificial neural networks: https://julienpascal.github.io/post/ann_2/

Graphical Processing Unit (GPU). In the end, the choice of method is up to the researcher, and should depend on the equipment available to her.

4.2 Stationary Equilibrium with Constant Aggregate State

We present two numerical methods to compute the constant aggregate state and associated stationary equilibrium in the model with econometric errors. We benchmark them against iterating over the social planner's value function W once it has been found by value function iteration. The main benefit of these two methods over value function iteration is speed and tractability, as will be made explicit below. We show that these methods can be adjusted in order to estimate the model parameters using data on who matched with whom. It is easy transition from computing a stationary equilibrium to estimating the model parameters while assuming that the data represent a stationary equilibrium. The minimal dataset comes from one market in stationary equilibrium and has cross sectional data on x, y, x', y' for matches x, y from that market, where x', y' are the states of the two matched agents at the beginning of the next period. Datasets with lengthier panels can also be used. In estimation, we assume that the transition rules for the individual worker and firm states P and Q are estimated in a first stage, as often done in dynamic discrete choice models (Rust (1987)). We focus on a second stage in which structural parameters are estimated. We first use the data on matches x, y to estimate the matching probabilities $\hat{\mu}_{xy}$ as well as $\hat{\mu}_{x0}$ and $\hat{\mu}_{0y}$ for unemployed workers and vacant jobs, respectively.

Mathematical Programming with Equilibrium Constraints

Mathematical Programming with Equilibrium Constraints, or MPEC, has been used by Su and Judd (2012) to estimate the dynamic discrete choice model of Rust (1987), and by Dubé et al. (2012) to estimate the aggregate demand model of Berry et al. (1995). MPEC formulates model solving or estimating parameters in a model as a constrained optimization problem, requiring nonlinear programming to numerically solve. Here, we apply MPEC to the two problems of computing a stationary equilibrium and estimating structural parameters.

To compute the stationary equilibrium, we solve the following set of equations for un-

knowns (U, V), (m, n):

$$
\begin{gathered}
\sum_{y \in \mathcal{Y}_0} \mu_{xy}(U,V,m,n) = m_x \quad \text{and} \quad \sum_{x \in \mathcal{X}_0} \mu_{xy}(U,V,m,n) = n_y, \\
\sum_{xy \in \mathcal{X}_0 \mathcal{Y}} P_{x'|xy} \mu_{xy}(U,V,m,n) = m_{x'} \quad \text{and} \quad \sum_{xy \in \mathcal{X} \mathcal{Y}_0} Q_{y'|xy} \mu_{xy}(U,V,m,n) = n_{y'}, \\
2 \sum_{xy \in \mathcal{X} \mathcal{Y}} \mu_{xy}(U,V,m,n) + \sum_{x \in \mathcal{X}} \mu_{x0}(U,V,m,n) + \sum_{y \in \mathcal{Y}} \mu_{0y}(U,V,m,n) = M.
\end{gathered}
\tag{3.36}
$$

Note that these are simply the feasibility conditions, stationary transition rules, and normalization equations outlined previously. In the model with Gumbel errors, by Proposition 1 the stationary matching policy is

$$
\begin{aligned}
\mu_{xy}(U,V,m,n) &= \sqrt{m_x n_y} \exp\left(\frac{\Phi_{xy} + \beta \sum_{x' \in \mathcal{X}} U_x P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} V_y Q_{y'|xy} - U_x - V_y}{2} \right) \\
\mu_{x0}(U,V,m,n) &= m_x \exp\left(\beta \sum_{x' \in \mathcal{X}} U_x P_{x'|xy} - U_x \right) \\
\mu_{0y}(U,V,m,n) &= n_y \exp\left(\beta \sum_{y' \in \mathcal{Y}} V_y Q_{y'|xy} - V_y \right)
\end{aligned}
\tag{3.37}
$$

To solve the system of equations (3.36) we use a nonlinear solver to maximize a constant function (say 0) subject to the constraints (3.36). We discuss in the next section our preferred choice of non-linear solver.

We now turn to estimation using MPEC. The functional form of the match production is parametrized by $\Phi_{xy} = \sum_k \lambda_k \phi_{xy}^k$ where ϕ_{xy}^k is the k^{th} basis function and the vector $\lambda = (\lambda)_k$ contains the coefficients on the basis function. The log likelihood for the data is

$$
2 \sum_{xy \in \mathcal{X} \mathcal{Y}} \hat{\mu}_{xy} \log \mu_{xy}(\lambda, U,V,m,n) + \sum_{x \in \mathcal{X}} \hat{\mu}_{x0} \log \mu_{x0}(\lambda, U,V,m,n) + \sum_{y \in \mathcal{Y}} \hat{\mu}_{0y} \log \mu_{0y}(\lambda, U,V,m,n),
\tag{3.38}
$$

where again $\hat{\mu}$ is the observed matching in the data. We maximize the log likelihood (3.38) over variables (U, V), (m, n) and λ subject to the constraints (3.36). This obtains the maximum likelihood estimation for λ, and is just a simple modification (adding an objective function) of the original program we have to compute the equilibrium in JuMP.

Reformulation as a min-max problem

Consider the following saddle-point problem

$$\max_{m,n} \min_{U,V} Z(U, V, U, V, m, n, \beta), \tag{3.39}$$

where

$$\begin{aligned} Z(U, V, U', V', m, n, \beta) = &- \sum_{x \in X} m_x - \sum_{y \in Y} n_y \\ &+ 2 \sum_{xy \in \mathcal{X}\,\mathcal{Y}} \mu_{xy}(U, V, U', V', m, n, \beta) \\ &+ \sum_{x \in \mathcal{X}} \mu_{x0}(U, V, U', V', m, n, \beta) + \sum_{y \in \mathcal{Y}} \mu_{0y}(U, V, U', V', m, n, \beta), \end{aligned} \tag{3.40}$$

and μ is defined as in (1) under Gumbel assumption 1 with μ's dependence in β having been made explicit. Problem (3.39)'s first order conditions are

$$\begin{aligned} &\frac{\partial Z}{\partial U_x} = - \sum_{y \in \mathcal{Y}_0} \mu_{xy}(U, V, U', V', m, n, \beta), \quad \frac{\partial Z}{\partial U'_x} = \beta \sum_{xy \in \mathcal{X}_0\,\mathcal{Y}} P_{x'|xy} \mu_{xy}(U, V, U', V', m, n, \beta) = 0, \\ &\frac{\partial Z}{\partial V_y} = - \sum_{x \in \mathcal{X}_0} \mu_{xy}(U, V, U', V', m, n, \beta), \quad \frac{\partial Z}{\partial V'_y} = \beta \sum_{xy \in \mathcal{X}\,\mathcal{Y}_0} Q_{y'|xy} \mu_{xy}(U, V, U', V', m, n, \beta) = 0, \\ &\frac{\partial Z}{\partial m_x} = -1 + \frac{1}{m_x} \sum_{y \in \mathcal{Y}_0} \mu_{xy}(U, V, U', V', m, n, \beta) = 0, \\ &\frac{\partial Z}{\partial n_y} = -1 + \frac{1}{n_y} \sum_{x \in \mathcal{X}_0} \mu_{xy}(U, V, U', V', m, n, \beta) = 0. \end{aligned} \tag{3.41}$$

When β is 1, these are exactly transition rules and feasibility conditions for aggregate steady state. One only has to solve problem (3.39) to find the stationary equilibrium and aggregate steady state. There exists easily implementable methods to solve min-max problem such as the Chambolle-Pock algorithm (Chambolle and Pock (2011)). However because in general $\beta < 1$, we cannot straightforwardly solve (3.39), as it is not the optimization problem whose first order conditions give us the stationary equilibrium: in $\frac{\partial Z}{\partial U_x}$ and $\frac{\partial Z}{\partial V_y}$ a β appears that takes us away from the stationary transition rules. To tackle this problem, we adapt the Chambolle-Pock algorithm to the case when $\beta < 1$ in order to cancel the extra β that appears in the first order conditions (3.41). To solve for (U, V), (m, n) that satisfy the

stationary equilibrium equations we choose step ϵ and we do

$$\begin{cases} \tilde{m}^{t+1} = 2m^t - m^{t-1} \\ \tilde{n}^{t+1} = 2n^t - n^{t-1} \\ U^{t+1} = U^t - \epsilon \left(\partial_U Z \left(U^t, V^t, U^t, V^t, \tilde{m}^{t+1}, \tilde{n}^{t+1}, \beta \right) \right) + \beta^{-1} \partial_{U'} Z \left(U^t, V^t, U^t, V^t, \tilde{m}^{t+1}, \tilde{n}^{t+1}, \beta \right) \\ V^{t+1} = V^t - \epsilon \left(\partial_V Z \left(U^t, V^t, U^t, V^t, \tilde{m}^{t+1}, \tilde{n}^{t+1}, \beta \right) \right) + \beta^{-1} \partial_{V'} Z \left(U^t, V^t, U^t, V^t, \tilde{m}^{t+1}, \tilde{n}^{t+1}, \beta \right) \\ m^{t+1} = m^t + \epsilon \partial_m Z \left(U^{t+1}, V^t, U^{t+1}, V^{t+1}, m^t, n^t, \beta \right) \\ n^{t+1} = n^{t+1} + \epsilon \partial_n Z \left(U^{t+1}, V^{t+1}, U^{t+1}, V^{t+1}, m^t, n^t, \beta \right). \end{cases} \tag{3.42}$$

Note that the algorithm is similar to an alternate gradient descent, which updates alternatively the parameters on which to minimize and those on which to maximize. The only difference is that instead of updating minimization parameters using the actual value of maximization paramters, we take an average between their actual and previous value. This simple adaptation is what ensures convergence of the algorithm in the case where $\beta = 1$, see Chambolle and Pock (2011) for more details. Note also that total mass of agents can be normalized through an additional equation, which simply adds a parameters on which to update the algorithm.

The min-max reformulation is easily adaptable to estimation. As for MPEC, assume the match production's functional form of is $\Phi_{xy}(\lambda) = \sum_k \lambda_k \phi_{xy}^k$ and $\hat{\mu}$ is the observed matching in the data. Then we can estimate $(\lambda_k)_k$ using moment condition

$$\sum_{xy \in \mathcal{X}_0 \mathcal{Y}_0} \Phi_{xy}^k \mu_{xy} = \sum_{xy \in \mathcal{X}_0 \mathcal{Y}_0} \Phi_{xy}^k \hat{\mu}_{xy} \quad \forall k. \tag{3.43}$$

To include this condition in the min-max formulation, simply write a new function $\hat{Z}$:

$$\begin{aligned} \hat{Z}(U, V, U', V', m, n, \beta, \lambda) = & -\sum_{x \in X} m_x - \sum_{y \in Y} n_y \\ & + 2 \sum_{xy \in \mathcal{X}\mathcal{Y}} \mu_{xy}(U, V, U', V', m, n, \beta, \lambda) \\ & + \sum_{x \in \mathcal{X}} \mu_{x0}(U, V, U', V', m, n, \beta, \lambda) \\ & + \sum_{y \in \mathcal{Y}} \mu_{0y}(U, V, U', V', m, n, \beta, \lambda) \\ & - \sum_{xy \in \mathcal{X}_0 \mathcal{Y}_0} \Phi_{xy}(\lambda) \hat{\mu}_{xy}(U, V, U', V', m, n, \beta, \lambda). \end{aligned} \tag{3.44}$$

Then first order conditions to the saddle-point problem

$$\max_{m,n} \min_{U,V,\lambda} \hat{Z}(U, V, U, V, m, n, \beta, \lambda) \tag{3.45}$$

are the same as (3.41), plus moment conditions (3.43). One only needs to adapt the algorithm in (3.42) to include the moment condition to estimate λ.

Comparison

We now proceed with some speed and precision comparisons between the two methods, MPEC and min-max. Speed is of course heavily dependent of the machine used, as well as of the solver chosen to solve the MPEC set of equations. The following benchmark are run using a computer with a 6-Core Intel processor, and 32GB of RAM. We define the MPEC optimization problem in the problem definition language JuMP (a Julia package) and solve the system using the nonlinear programming solver MadNLP[3] based on Shin et al. (2020) which relies on interior point methods to solve non-linear optimization problems. Note that much of the speed gains from MPEC rest on the use of the programming solver: for instance the Ipopt solver performs remarkably less well than MadNLP, although it is often used to solve non linear problems and relies on the same interior point methods as MadNLP. On the other hand, the min-max formulation and adapted Chambolle-Pock algorithm are more robust to implementation choices, as they can easily be coded from scratch by the researcher in any computing language of her choice.

Table 3.1 and 3.2 reports both methods' speed of convergence for equilibrium computation (assuming match surplus is known) and estimation depending either on the number of types of agents on each side of the market (for equilibrium computation) or on the number of surplus parameters to estimate (for estimation). Both methods are extremely fast when the number of types or parameters is small. They both slow down as numbers grow, although the min-max method remains fast.

Table 3.1: Equilibrium computation: methods' speed depending on the number of types

Method	$nbx = 2$, $nby = 2$	$nbx = 10$, $nby = 10$	$nbx = 30$, $nby = 30$
MPEC	0.004s	1.168s	408.770s
MinMax	0.048s	1.805s	35.871s

Average speed in seconds over 10 runs

[3]Downloadable from Github: https://github.com/MadNLP/MadNLP.jl

Table 3.2: Estimation: methods' speed depending on the number of parameters

Method	$nbk = 2$	$nbk = 10$	$nbk = 30$
MPEC	3.077s	11.918s	35.910s
MinMax	0.751s	2.044s	9.211s

Average speed in seconds over 10 runs
$nbx = 10$, $nby = 10$

Table reports the largest standard error obtained with each method depending on how big the sample from which $\hat{\mu}$ is obtained is. Both methods yields satisfactory standard errors when the sample size is large.

Table 3.3: Estimation: methods' estimation precision depending on sample size

Method	$obs = 1e4$	$obs = 1e5$	$obs = 1e6$
MPEC	0.215	0.062	0.020
MinMax	0.216	0.063	0.020

Largest standard error over 10 runs
$nbx = 10$, $nby = 10$, $nbk = 10$

5 Empirical Application

To illustrate the usefulness of our model applied to labor data, we estimate moving costs for Swedish engineers in the 1970s by age and geographic distance. We show costs to switching region for work are sizeable, and non linear in age.

The dataset used covers a subset of Swedish engineers and the firms they work for from 1970 to 1990 (see Fox (2009), Fox (2010a) for more background on the data). Observations are at individual times year level. Workers and firms are observed each year with a unique identifier. The data contains a number of characteristics on both workers and firms, among which worker age and firm location. Worker and firm type are defined as follows:

$$x = \{\text{age}, \text{previous location}\} \text{ and } y = \{\text{location}\},$$

where the previous location for a type x workers is their employer's location in year $t-1$, which is assumed to become the worker's location in year t, and location for type y firms is its geographic region in year t. Both ages and locations are aggregated from the data: we

assume 5 different ages in the model: up to 25 years old (age of 1), from 26 to 35 (age of 2), from 36 to 45 (age of 3), from 46 to 55 (age of 4), 56 years old and more (age fo 5). Location is provided in the raw data as one of Sweden's 23 counties. We aggregate these into larger areas based on county location to obtain four regions: Stockholm, counties immediately next to Stockholm, counties in south Sweden, and counties in the center and north of the country.

As a result, there are 20 types for workers and 4 types for firms: $nbx = 20$ and $nby = 4$. Following Fox (2010a), we are interested in evaluating switching costs between regions by age. Our match production function is

$$\Phi(x, y) = \sum_{a=1}^{5} \lambda_a \mathbb{1}_{[x_{age}=a]} \text{dist}(x, y),$$

where $\text{dist}(x, y)$ is the average distance between counties that make up the worker's previous region in x and the firm's region in y. We are looking to estimate parameters λ_a for each age bin a.

Let $\left(\tilde{\mu}^t_{xy}\right)_{x,y}$ be the observed matching in year t, i.e. $\tilde{\mu}^t_{xy}$ is the empirical probability of randomly drawing a match between a worker of type x and a firm of type y in the observed data. Set $\tilde{\mu}_{xy} = \sum_{t=1971}^{1980} \frac{1}{10} \tilde{\mu}^t_{xy}$ to be the average of this probability over the decade. We do not observe unemployed workers nor firms with no employee and hence include no unmatched agents in the estimation.

Transition matrices are estimated directly from the data, using $\tilde{\mu}$. Workers transition deterministically from their location at $t-1$ to their employers' location at t, so that $x^t_{loc} = y^{t-1}_{loc}$. We assume a probability $\rho = .1$ of ageing from age a to age $a+1$. Engineers regularly leave the market either to retire or work in industries that are not accounted for in the data, and a number of firms also leaves the market every year. Therefore we assume some attrition in both populations. We compute attrition rates δ_x and δ_y by worker and firm type. They are taken to be the average of the share of agent of each type who leave every year between 1971 and 1980. Worker transition matrix P is a 20×80 matrix with entries

$$P_{x'|xy} = \begin{cases} (1-\rho)\delta_x \tilde{\mu}_{xy} \text{ if } x_{age} = x'_{age} \text{ and } y_{loc} = x'_{loc} \\ \rho \delta_x \tilde{\mu}_{xy} \text{ if } x_{age} + 1 = x'_{age} \text{ and } y_{loc} = x'_{loc} \\ 0 \text{ otherwise.} \end{cases}$$

Firm type transition matrix Q is a 4×80 matrix with entries

$$Q_{y'|xy} = \begin{cases} \delta_y \tilde{\mu}_{xy} \text{ if } y_{loc} = y'_{loc} \\ 0 \text{ otherwise.} \end{cases}$$

In order to remain on the stationary framework despite workers' and firms' attrition, we introduce an incoming flow $(i_x)_x$ for workers and $(i_y)_y$ that is computed to compensate loss from attrition based on observed matching:

$$i_{x'} = \sum_y \tilde{\mu}_{x'y} - \sum_{x,y} P_{x'|xy} \tilde{\mu}_{xy} \text{ and } i_{y'} = \sum_x \tilde{\mu}_{xy'} - \sum_{x,y} Q_{y'|xy} \tilde{\mu}_{xy}. \tag{3.46}$$

Accounting for attrition and incoming flows slightly modifies the stationary equilibrium equations that become:

$$\begin{gathered} \sum_{y \in \mathcal{Y}_0} \mu_{xy}(U, V, m, n) = m_x \quad \text{and} \quad \sum_{x \in \mathcal{X}_0} \mu_{xy}(U, V, m, n) = n_y, \\ \sum_{xy \in \mathcal{X}_0 \mathcal{Y}} P_{x'|xy} \mu_{xy}(U, V, m, n) = m_{x'} + i_{x'} \quad \text{and} \quad \sum_{xy \in \mathcal{X} \mathcal{Y}_0} Q_{y'|xy} \mu_{xy}(U, V, m, n) = n_{y'} + i_{y'}, \\ 2 \sum_{xy \in \mathcal{X} \mathcal{Y}} \mu_{xy}(U, V, m, n) + \sum_{x \in \mathcal{X}} \mu_{x0}(U, V, m, n) + \sum_{y \in \mathcal{Y}} \mu_{0y}(U, V, m, n) = M, \end{gathered} \tag{3.47}$$

where μ is unchanged as in (3.37).

To complete the estimation, we assume discount factor $\beta = .95$. Point estimates obtained with Chambolle-Pock and MPEC are as follows:

Table 3.4: Estimates for moving cost by age bin

	λ_1	λ_2	λ_3	λ_4	λ_5
Chambolle-Pock	-51.81	-49.86	-49.10	-47.28	-52.97
MPEC	-48.42	-48.88	-51.22	-48.31	-50.32

Estimates are non linear in age: moving costs decrease as workers become older, except for 56 years old and more who face higher moving costs with respect than their younger peers. Because unemployed workers and vacant firms are not accounted for in this estimation, match production is only identified up to a constant. Thus estimates are only interpretable relative to each other: for instance, moving when a worker is between 26 and 35 years old imposes an extra penalty of about 2 on production per extra kilometer compared to when a worker is up to 25 years old. To interpret these results, we should also keep in mind that switching costs

apply to total match surplus, which accrues both to firm production and worker amenity. Hence switching costs can be due either to a immediate loss in match production for the firm, or to a dislike for moving to a different region from workers. It is likely to two combine to different degrees depending on age, which might help explain the non-linearities in switching costs by age.

6 Conclusion

In this chapter, we develop a model of dynamic matching with a wide range of applications, to labor, family economics and industrial organization. The main features of the model are that 1) the matching game is repeated every period and 2) agents are forward-looking and account in changes in their state variable caused by who they match with in the present period. We expose the model without and with econometric shocks, and show equilibrium can be found in both cases by solving the social planner Bellman question. Importantly, we show that a stationary equilibrium, which does not change the distribution of agents' state variables in every period, always exists both without and with econometric shocks. Then, we provide methods to compute both the dynamic and the stationary equilibrium: equilibrium can always be found in general by value function iteration on the social planner Bellman equation, or through deep learning methods that fit artificial neural networks to minimize a loss function that approximate the social planner Bellman equation's fixed point. Then, we introduce two methods to compute the stationary equilibrium: by solving for the stationary equilibrium equations using a non-linear solver, or by reformulating the equations as a min-max optimization problem. We compare the two methods in terms of speed and precision, and show that the min-max reformulation is faster when the number of types is large, although the two methods have very similar precision. Finally, we adapt our methods to estimation and apply them to estimating geographic mobility costs for Swedish engineers in the 1970s. We find the mobility costs are substantial and non-linear in age.

References

Barany, Z. and Holzheu, K. (2021). Worker Specialization and the Aggregate Economy. *Working Paper*.

Becker, G. S. (1973). A Theory of Marriage: Part I. *Journal of Political Economy*, 81(4):813–846.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.

Berry, S. T. and Compiani, G. (2020). An Instrumental Variable Approach to Dynamic Models. NBER Working Paper 27756, National Bureau of Economic Research, Inc.

Burdett, K. and Mortensen, D. T. (1998). Wage Differentials, Employer Size, and Unemployment. *International Economic Review*, 39(2):257–273.

Chambolle, A. and Pock, T. (2011). A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145.

Chiappori, P.-A., Salanié, B., and Weiss, Y. (2017). Partner Choice, Investment in Children, and the Marital College Premium. *American Economic Review*, 107(8):2109–2167.

Choo, E. (2015). Dynamic Marriage Matching: An Empirical Framework. *Econometrica*, 83(4):1373–1423.

Choo, E. and Siow, A. (2006). Who Marries Whom and Why. *Journal of Political Economy*, 114(1):175–201.

Dagsvik, J. K. (2000). Aggregation in Matching Markets. *International Economic Review*, 41(1):27–57.

Dubé, J.-P., Fox, J. T., and Su, C.-L. (2012). Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation. *Econometrica*, 80(5):2231–2267.

Dupuy, A. and Galichon, A. (2014). Personality Traits and the Marriage Market. *Journal of Political Economy*, 122(6):1271–1319.

Erlinger, A., McCann, R. J., Shi, X., Siow, A., and Wolthoff, R. (2015). Academic wages and pyramid schemes: A mathematical model. *Journal of Functional Analysis*, 269(9):2709–2746.

Fang, H.-r. and Saad, Y. (2009). Two classes of multisecant methods for nonlinear acceleration. *Numerical Linear Algebra with Applications*, 16(3):197–221.

Fox, J. (2010a). Estimating the Employer Switching Costs and Wage Responses of Forward-Looking Engineers. *Journal of Labor Economics*, 28(2):357–412.

Fox, J., Yang, C., and Hsu, D. H. (2018). Unobserved Heterogeneity in Matching Games. *Journal of Political Economy*, Vol 126, No 4.

Fox, J. T. (2009). Firm-Size Wage Gaps, Job Responsibility, and Hierarchical Matching. *Journal of Labor Economics*, 27(1):83–126. Publisher: University of Chicago Press.

Fox, J. T. (2010b). Identification in matching games. *Quantitative Economics*, 1(2):203–254.

Fox, J. T. (2018). Estimating matching games with transfers. *Quantitative Economics*, 9(1):1–38.

Gale, D. (1989). *The Theory of Linear Economic Models*. University of Chicago Press, Chicago, IL.

Galichon, A. and Salanié, B. (2021). Cupid's Invisible Hand: Social Surplus and Identification in Matching Models. *The Review of Economic Studies*.

Gretsky, N., Ostroy, J., and Zame, W. (1992). The Nonatomic Assignment Model. *Economic Theory*, 2:103–27.

Koopmans, T. C. and Beckmann, M. (1957). Assignment Problems and the Location of Economic Activities. *Econometrica*, 25(1):53–76.

Krusell, P. and Smith, Anthony A., J. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Lazear, E. (2009). Firm-Specific Human Capital: A Skill-Weights Approach. *Journal of Political Economy*, 117(5):914–940.

McCann, R. J., Shi, X., Siow, A., and Wolthoff, R. (2015). Becker Meets Ricardo: Multisector Matching with Communication and Cognitive Skills. *Journal of Law, Economics, and Organization*, 31(4):690–720.

Miller, R. A. (1984). Job Matching and Occupational Choice. *Journal of Political Economy*, 92(6):1086–1120.

Pakes, A. (1986). Patents as Options: Some Estimates of the Value of Holding European Patent Stocks. *Econometrica*, 54(4):755–784.

Peski, M. (2021). Tractable Model of Dynamic Many-to-Many Matching. *American Economic Journal: Microeconomics*.

Rios-Rull, J.-V. (1995). Models with Heterogeneous Agents. In *Frontiers of Business Cycle Research*, pages 98–125. Princeton University Press.

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, 55(5):999–1033.

Shapley, L. S. and Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130.

Shin, S., Coffrin, C., Sundar, K., and Zavala, V. M. (2020). Graph-Based Modeling and Decomposition of Energy Infrastructures. *arXiv preprint arXiv:2010.02404*.

Su, C.-L. and Judd, K. L. (2012). Constrained Optimization Approaches to Estimation of Structural Models. *Econometrica*, 80(5):2213–2230.

Walker, H. and Ni, P. (2011). Anderson Acceleration for Fixed-Point Iterations. *SIAM J. Numerical Analysis*, 49:1715–1735.

Weintraub, G. Y., Benkard, C. L., and Van Roy, B. (2008). Markov Perfect Industry Dynamics With Many Firms. *Econometrica*, 76(6):1375–1411.

Wolpin, K. I. (1984). An Estimable Dynamic Stochastic Model of Fertility and Child Mortality. *Journal of Political Economy*, 92(5):852–874.

A Proofs

Theorem 2.1

The central planner problem (3.7), (3.8), (3.9) has dual:

$$\min_{U_x, V_y} \left\{ \sum_{x \in \mathcal{X}} m_x^0 U_x^0 + \sum_{y \in \mathcal{Y}} n_y^0 V_y^0 \right\} \tag{3.48}$$

subject to stability conditions

$$\begin{aligned} U_x^t + V_y^t &\geq (\alpha_{xy} + \gamma_{xy}) + \beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1} P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1} Q_{y'|xy} \quad \forall\, t, x, y \\ U_x^t &\geq +\beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1} P_{x'|x0} \quad \forall\, t, x \\ V_y^t &\geq +\beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1} Q_{y'|0y} \quad \forall\, t, y. \end{aligned} \tag{3.49}$$

Let $(\mu^t)_t$ be a solution to primal problem (3.7), (3.8), (3.9) and $(U^t, V^t)_t$ be a solution to dual problem (3.48), (3.49). Introduce $(w^t)_t$ such that

$$U_x^t - \alpha_{xy} - \beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1} P_{x'|xy} \geq w_{xy}^t \geq -V_y^t + \gamma_{xy}^t + \beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1} Q_{y'|xy}. \tag{3.50}$$

Then $(\mu^t, w^t)_t$ satisfies (3.6) and is therefore a competitive equilibrium.

Conversely, assume $(\mu^t, w^t)_t$ be a competitive equilibrium, i.e. it satisfies (3.6) for $(U^t, V^t)_t$ defined through (3.3) and (3.5). Then wage $(\mu^t, w^t)_t$ satisfy primal feasibility

$$\begin{cases} \mu_{xy}^t \geq 0 \quad \forall t, x, y \\ \sum_{y \in \mathcal{Y}_0} \mu_{xy}^t = m_x^t \quad \forall\, t, x \quad \text{and} \quad \sum_{x \in \mathcal{X}_0} \mu_{xy}^t = n_y^t \quad \forall\, t, y \\ \sum_{x'y' \in \mathcal{X}\,\mathcal{Y}_0} P_{x|x'y'} \mu_{x'y'}^t = m_x^{t+1} \quad \forall\, t, x \quad \text{and} \quad \sum_{x'y' \in \mathcal{X}_0\,\mathcal{Y}} Q_{y|x'y'} \mu_{xy}^t = m_x^{t+1} \quad \forall\, t, y, \end{cases} \tag{3.51}$$

dual feasibility

$$\begin{cases} U_x^t + V_y^t \geq (\alpha_{xy} + \gamma_{xy}) + \beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1} P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1} Q_{y'|xy} \quad \forall\, t, x, y \\ U_x^t \geq +\beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1} P_{x'|x0} \quad \forall\, t, x \\ V_y^t \geq +\beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1} Q_{y'|0y} \quad \forall\, t, y, \end{cases} \tag{3.52}$$

and complementary slackness

$$\begin{cases} \mu_{xy}^t \left(U_x^t + V_y^t - (\alpha_{xy} + \gamma_{xy}) - \beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1} P_{x'|xy} - \beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1} Q_{y'|xy} \right) = 0 \quad \forall t, x, y \\ \mu_{x0}^t \left(U_x^t - \beta \sum_{x' \in \mathcal{X}} U_{x'}^{t+1} P_{x'|x0} \right) = 0 \quad \forall t, x \\ \mu_{0y}^t \left(V_y^t - \beta \sum_{y' \in \mathcal{Y}} V_{y'}^{t+1} Q_{y'|0y} \right) = 0 \quad \forall t, y, \end{cases} \tag{3.53}$$

for the social planner problem and its dual.

Theorem 2.2

We show that the map $\varphi(W) \to \max_{\mu \in \mathcal{M}(m,n)} \left\{ \sum_{xy \in \mathcal{X}_0 \mathcal{Y}_0} \mu_{xy} \left(\alpha_{xy} + \gamma_{xy} \right) + \beta W \left(P\mu, Q\mu \right) \right\}$ is a contraction, where $\mathcal{M}(m,n) = \left\{ \mu_{xy} \geq 0 | \sum_{y \in \mathcal{Y}_0} \mu_{xy} = m_x, \sum_{x \in \mathcal{X}_0} \mu_{xy} = n_y \, \forall \, x, y \right\}$.

Fix (m,n), consider two different continuation values W, W' and define

$$\tilde{\mu} \in \underset{\mu \in \mathcal{M}(m,n)}{\arg\max} \left\{ \sum_{xy \in \mathcal{X}_0 \mathcal{Y}_0} \mu_{xy} \left(\alpha_{xy} + \gamma_{xy} \right) + \beta W \left(P\mu, Q\mu \right) \right\} \tag{3.54}$$

then

$$\varphi(W) - \varphi(W') \leq \beta \left(W(\tilde{m}, \tilde{n}) - W'(\tilde{m}, \tilde{n}) \right) \tag{3.55}$$

where $(\tilde{m}, \tilde{n}) = (P\tilde{\mu}, Q\tilde{\mu})$. Hence

$$\varphi(W) - \varphi(W') \leq \beta \max_{\tilde{m}, \tilde{n}} \left(W(\tilde{m}, \tilde{n}) - W'(\tilde{m}, \tilde{n}) \right) \tag{3.56}$$

where $\max_{\tilde{m}, \tilde{n}}$ is taken over $(\tilde{m}, \tilde{n})$ such that $\sum_x \tilde{m}_x = \sum_x m_x$ and $\sum_y \tilde{n}_y = \sum_y n_y$. Because the above equation is true for any (m,n) we have

$$\|\varphi(W) - \varphi(W')\| \leq \beta \|W - W'\| \tag{3.57}$$

where $\|.\|$ is the sup-norm. Therefore φ us a contraction for the sup-norm, and a fixed point of ϕ exists and is unique.

Theorem 2.3

A competitive equilibrium (μ, w) is defined through equations:

$$\begin{cases} \mu_{xy} \geq 0, \quad \mu_{x0} \geq 0, \quad \mu_{0y} \geq 0 \quad \forall\, x, y \\ U_x + V_y \geq (\alpha_{xy} + \gamma_{xy}) + \beta \sum_{x' \in \mathcal{X}} U_{x'} P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} V_{y'} Q_{y'|xy} \quad \forall\, x, y \\ U_x \geq \beta \sum_{x' \in \mathcal{X}} U_{x'} P_{x'|x0} \quad \forall\, x \\ V_y \geq \beta \sum_{y' \in \mathcal{Y}} V_{y'} Q_{y'|0y} \quad \forall\, y \\ \mu_{xy} > 0 \Rightarrow U_x + V_y - (\alpha_{xy} + \gamma_{xy}) - \beta \sum_{x' \in \mathcal{X}} U_{x'} P_{x'|xy} - \beta \sum_{y' \in \mathcal{Y}} V_{y'} Q_{y'|xy} = 0 \quad \forall\, x, y \\ \mu_{x0} > 0 \Rightarrow U_x - \beta \sum_{x' \in \mathcal{X}} U_{x'} P_{x'|x0} = 0 \quad \forall\, x \\ \mu_{0y} > 0 \Rightarrow V_y - \beta \sum_{y' \in \mathcal{Y}} V_{y'} Q_{y'|0y} \;\; = 0 \quad \forall\, y. \end{cases} \tag{3.58}$$

(μ, w) is stationary iff

$$\sum_{y \in \mathcal{Y}_0} \mu_{xy} = \sum_{x'y' \in \mathcal{X}\,\mathcal{Y}_0} P_{x|x'y'} \mu_{x'y'} \quad \forall\, x \quad \text{and} \quad \sum_{x \in \mathcal{X}_0} \mu_{xy} = \sum_{x'y' \in \mathcal{X}_0\,\mathcal{Y}} Q_{y|x'y'} \mu_{x'y'} \quad \forall\, y. \tag{3.59}$$

Normalize $M = 1$ without loss of generality. The competitive μ also satisfies

$$2 \sum_{xy \in \mathcal{X}\,\mathcal{Y}} \mu_{xy} + \sum_{x \in \mathcal{X}} \mu_{x0} + \sum_{y \in \mathcal{Y}} \mu_{0y} = 1. \tag{3.60}$$

Now introduce vector $\delta = (\delta_x, \delta_y)_{x,y}$, such that $\delta \geq 0$ and $\sum_x \delta_x + \sum_y \delta_y = \frac{1}{1-\beta}$. Let us show that there exists a μ, and (U, V) that satisfy (3.58), (3.60) and a relaxed version of (3.59):

$$\sum_{y \in \mathcal{Y}_0} \mu_{xy} - \beta \sum_{x'y' \in \mathcal{X}\,\mathcal{Y}_0} P_{x|x'y'} \mu_{x'y'} = \delta_x \quad \text{and} \quad \sum_{x \in \mathcal{X}_0} \mu_{xy} - \beta \sum_{x'y' \in \mathcal{X}_0\,\mathcal{Y}} Q_{y|x'y'} \mu_{x'y'} = \delta_y. \tag{3.61}$$

Equations (3.58), (3.60), (3.61) are the optimality conditions for the following linear programming problem:

$$\begin{gathered} \max_{\mu \geq 0} \sum_{xy \in \mathcal{X}\,\mathcal{Y}} \Phi_{xy} \mu_{xy} \\ \text{s.t (3.61).} \end{gathered} \tag{3.62}$$

A solution μ and associated Lagrange multipliers (U, V) exist to problem (3.62) for any vector δ.

We now look for δ such that the optimum $\mu(\delta)$ satisfies

$$(1-\beta) \sum_{x'y' \in \mathcal{X}\,\mathcal{Y}_0} P_{x|x'y'} \mu_{x'y'}(\delta) = \delta_x \quad \forall\, x \quad \text{and} \quad (1-\beta) \sum_{x'y' \in \mathcal{X}_0\,\mathcal{Y}} Q_{y|x'y'} \mu_{x'y'}(\delta) = \delta_y \quad \forall\, y. \tag{3.63}$$

Then the solution $\mu(\delta)$ to (3.62) satisfies stationarity conditions (3.59) and normalization (3.60).

Let φ be the map such that $\varphi_x(\delta) = (1-\beta)\sum_{x'y' \in \mathcal{X}\mathcal{Y}_0} P_{x|x'y'}\mu_{x'y'}(\delta)$ and $\varphi_y(\delta) = (1-\beta)\sum_{x'y' \in \mathcal{X}_0\mathcal{Y}} Q_{y|x'y'}\mu_{x'y'}(\delta)$. In short $\varphi(\delta) = (1-\beta)(P\mu(\delta), Q\mu(\delta))$. We are looking for a fixed point of φ. Let us show that φ has closed graph and that $\varphi(\delta)$ is non-empty and convex to apply Kakutani theorem.

The graph if φ is $G_\varphi = \left\{(b, (1-\beta)(P\mu, Q\mu)) \text{ s.t } \sum_{xy\in\mathcal{X}\mathcal{Y}} \Phi_{xy}\mu_{xy} = f(b)\right\}$, where f is defined through

$$f(b) = \max_{\mu \geq 0} \left\{ \sum_{xy \in \mathcal{X}\mathcal{Y}} \Phi_{xy}\mu_{xy} \text{ s.t } (3.61) \right\}. \tag{3.64}$$

G_φ is closed by continuity of function f. Next, $\varphi(\delta)$ is non-empty because there always exist a solution to (3.62). Finally $\varphi(\delta)$ is convex: let $(1-\beta)(P\mu(\delta), Q\mu(\delta))$, $(1-\beta)(P\mu'(\delta), Q\mu'(\delta)) \in \varphi(\delta)$, with $\mu \neq \mu'$. Then $(1-t)(1-\beta)(P\mu(\delta), Q\mu(\delta)) + t(1-\beta)(P\mu'(\delta), Q\mu'(\delta)) \in \varphi(\delta)$.

By Kakutani theorem, φ admits a fixed point, and there exists δ that satisfies (3.63).

Theorem 3.1

We follow Gretsky et al. (1992) and Galichon and Salanié (2021). Consider the following problem of minimizing the sum of individual welfare under stability conditions:

$$\begin{aligned} \min_{(u^t)_t, (v^t)_t} & \sum_i u_i^1 + \sum_j v_j^1 \\ \text{s.t } & u_i^t + v_j^t \geq \Phi_{x_i y_j} + \epsilon_{iy} + \eta_{xj} \\ & \qquad + \beta\mathbb{E}_P\left[u^{t+1}|x_i y_j\right] + \beta\mathbb{E}_Q\left[v^{t+1}|x_i y_j\right] \quad \forall\, t, x, y \\ & u_i^t \geq + \epsilon_{i0} + \beta\mathbb{E}_P\left[u^{t+1}|x_i 0\right] \quad \forall\, t, x \\ & v_j^t \geq \eta_{0j} + \beta\mathbb{E}_Q\left[v^{t+1}|0 y_j\right] \quad \forall\, t, y. \end{aligned} \tag{3.65}$$

Take any two $(u^t)_t$ and $(v^t)_t$ such that $u^t_{xy} + v^t_{xy} \geq \Phi_{xy} + \beta\mathbb{E}_P\left[\max_{y\in\mathcal{Y}_0}\left\{u^{t+1}_{xy} + \epsilon_y\right\}|xy\right] + \beta\mathbb{E}_Q\left[\max_{x\in\mathcal{X}_0}\left\{v^{t+1}_{xy} + \eta_x\right\}|xy\right]$, and $u_{x0} = 0$, $u_{y0} = 0$. Define

$$u_i^t = \max_y\left\{u^t_{xy} + \epsilon_{iy}\right\} \quad \forall\, t, x \quad \text{and} \quad v_j^t = \max_y\left\{v^t_{xy} + \eta_{xj}\right\} \quad \forall\, t, y. \tag{3.66}$$

Then $(u_i^t, v_j^t)_t$ satisfy problem (3.65)'s constraints. Reciprocally, fix any $(u_i^t, v_j^t)_t$ that satisfy

the constraints in this problem and let

$$u_{xy}^t = \min_{i, x_i = x} \left\{ u_i^t - \epsilon_{iy} \right\}, \quad u_{x0}^t = 0 \quad \text{and} \quad v_{xy}^t = \min_{j, y_j = y} \left\{ v_j^t - \eta_{xj} \right\}, \quad v_{0y}^t = 0. \tag{3.67}$$

Then the constraint in problem (3.65) becomes $u_{xy}^t + v_{xy}^t \geq \Phi_{xy} + \beta \mathbb{E}_P \left[\max_{y \in \mathcal{Y}_0} \left\{ u_{xy}^{t+1} + \epsilon_y \right\} |xy \right] + \beta \mathbb{E}_Q \left[\max_{x \in \mathcal{X}_0} \left\{ v_{xy}^{t+1} + \eta_x \right\} |xy \right]$.

Applying the law of large numbers, we obtain that problem (3.65) is equivalent to

$$\begin{aligned} \min_{u,v} & \sum_{x \in \mathcal{X}} m_x^1 G_x(u^1) + \sum_{y \in \mathcal{Y}} n_y^1 H_y(v^1) \\ & \text{s.t } u_{xy}^t + v_{xy}^t \geq \Phi_{xy} + \beta \mathbb{E}_P \left[G_x(u^{t+1}) | xy \right] + \beta \mathbb{E}_Q \left[H_y(v^{t+1}) | xy \right] \quad \forall\, t, x, y \\ & \quad u_{x0}^t \geq \beta \mathbb{E}_P \left[G_x(u^{t+1}) | x0 \right] \quad \forall\, t, x \\ & \quad v_{0y}^t \geq \beta \mathbb{E}_Q \left[H_y(v^{t+1}) | 0y \right] \quad \forall\, t, y, \end{aligned}$$

which is the dual (3.26) of social planner problem (3.20). To see this, rewrite (3.26) as a saddle point problem, and add the term

$$\sum_{t=1} \left(\sum_x m_x^t G_x(u^t) + \sum_y n_y^t H_y(v^t) \right) - \sum_{t=0} \left(\sum_x m_x^{t+1} G_x(u^{t+1}) + \sum_y n_y^{t+1} H_y(v^{t+1}) \right).$$

Theorem 3.2
A stationary equilibrium's matching policy under Gumbel shocks writes

$$\begin{aligned} \mu_{xy}(U, V, m, n) &= \sqrt{m_x n_y} \exp \left(\frac{\Phi_{xy} + \beta \sum_{x' \in \mathcal{X}} U_x P_{x'|xy} + \beta \sum_{y' \in \mathcal{Y}} V_y Q_{y'|xy} - U_x - V_y}{2} \right) \\ \mu_{x0}(U, V, m, n) &= m_x \exp \left(\beta \sum_{x' \in \mathcal{X}} U_x P_{x'|xy} - U_x \right) \\ \mu_{0y}(U, V, m, n) &= n_y \exp \left(\beta \sum_{y' \in \mathcal{Y}} V_y Q_{y'|xy} - V_y \right), \end{aligned} \tag{3.68}$$

where the dependence of matching policy μ on expected indirect payoffs (U, V) and margins (m, n) has been made explicit. To show that a stationary equilibrium exists, we must show that there exists a tuple (U, V, m, n) that satisfies feasibility constraint:

$$\sum_{y \in \mathcal{Y}_0} \mu_{xy}(U, V, m, n) = m_x \quad \text{and} \quad \sum_{x \in \mathcal{X}_0} \mu_{xy}(U, V, m, n) = n_y, \tag{3.69}$$

and transition rules

$$\sum_{xy\in\mathcal{X}_0\,\mathcal{Y}} P_{x'|xy}\mu_{xy}\left(U,V,m,n\right)=m_{x'} \quad \text{and} \quad \sum_{xy\in\mathcal{X}\,\mathcal{Y}_0} Q_{y'|xy}\mu_{xy}\left(U,V,m,n\right)=n_{y'}. \tag{3.70}$$

We proceed as follows: in a first step, we show that given (m,n) the map $F_{m,n}:(U,V)\to(U,V)$ defined through feasibility constraints is a contraction for the sup norm, and hence admits a fixed point $(\bar{U}(\bar{n},\bar{m}),\bar{V}(\bar{n},\bar{m})$. In a second step, we show that given (U,V), the map $G:(m,n)\to(m,n)$ is continuous on a convex compact. Hence by Brouwer's theorem it admits a fixed point $(\bar{n},\bar{m})$. The tuple $(\bar{U}(\bar{n},\bar{m}),\bar{V}(\bar{n},\bar{m}),\bar{n},\bar{m})$ verifies the conditions for the stationary equilibrium.

Step 1: Define map $F_{m,n}:(U,V)\to(U,V)$ as $F_{m,n}=F^1\circ F^2_{n,m}$ with

$$\begin{aligned}
\left(F^1(U,V)\right)_{xy}&=\frac{\Phi_{xy}+\beta\left(PU+QV\right)_{xy}}{2}\\
\left(F^1(U,V)\right)_{x0}&=\beta\left(PU\right)_{x0}\\
\left(F^1(U,V)\right)_{0y}&=\beta\left(QV\right)_{0y},
\end{aligned} \tag{3.71}$$

where $(PU)_{xy}$ is short for $\sum_{x'}P_{x'|xy}U_{x'}$. F^1 is continuous in (U,V). Map $F^2_{n,m}:\Phi\to(U,V)$ is implicitly defined through equations

$$\begin{aligned}
&\sum_{y\in\mathcal{Y}}\sqrt{m_xn_y}\exp\left(\frac{\Phi_{xy}-U_x-V_y}{2}\right)+m_x\exp\left(\Phi_{x0}-U_x\right)=m_x\\
&\sum_{x\in\mathcal{X}}\sqrt{m_xn_y}\exp\left(\frac{\Phi_{xy}-U_x-V_y}{2}\right)+n_y\exp\left(\Phi_{0y}-V_y\right)=n_y.
\end{aligned} \tag{3.72}$$

Let $S=\begin{pmatrix}\frac{1}{2}I_X\otimes 1_Y & \frac{1}{2}1_X\otimes I_Y\\ I_X & 0_{X\times Y}\\ 0_{Y\times X} & I_Y\end{pmatrix}$ where I_X is the identity matrix of size $X=\#\mathcal{X}$ and 1_X is a column vector of ones of size X. The same goes for I_Y and 1_Y of size $Y=\#Y$. Also let $\phi_{xy}=\Phi_{xy}+\log n_x+\log m_y$, $\Phi_{x0}=\log n_x$, $\Phi_{0y}=\log m_y$, then F_2 rewrites in matrix form:

$$S^\top\exp\left(-Sp+\phi\right)-q=0, \tag{3.73}$$

where $p=(U,V)$ and $q=(m,n)$.

We can apply the implicit function theorem to get

$$D_\phi p = \left[S^\top \exp\left(-Sp + \phi\right) S\right]^{-1} S^\top \exp\left(-Sp + \phi\right). \tag{3.74}$$

We now aim to show that $D_\phi p$ is bounded by 1. If $D_\phi p$ is bounded by 1, then p is a Lipschitz continuous function of ϕ with Lipschitz constant 1. Then $F_{m,n}$ is a Lipschitz continuous map for any given m, n with Lipschitz constant β.

To show that $D_\phi p$ is bounded, take w some vector of size $X \times Y + X + Y$ with $\|w\| < k$ for some $k > 0$. Then

$$\begin{aligned}\left(S^\top \exp\left(-Sp + \phi\right) w\right)_x &\leq \frac{1}{2}\sum_{y\in\mathcal{Y}} \sqrt{n_x m_y} \exp\left(\frac{\phi_{xy} - U_x - V_y}{2}\right) k + n_x \exp\left(\phi_{x0} - U_x\right) k \quad \forall x \\ \left(S^\top \exp\left(-Sp + \phi\right) w\right)_y &\leq \frac{1}{2}\sum_{x\in\mathcal{X}} \sqrt{n_x m_y} \exp\left(\frac{\phi_{xy} - U_x - V_y}{2}\right) k + m_y \exp\left(\phi_{0y} - V_y\right) k \quad \forall y.\end{aligned} \tag{3.75}$$

Therefore

$$\begin{aligned}&\left[S^\top \exp\left(-Sp + \Phi\right) S\right]^{-1} S^\top \exp\left(-Sp + \Phi\right) w \leq [k, \dots, k]^\top \\ \Rightarrow &\left\|\left[S^\top \exp\left(-Sp + \Phi\right) S\right]^{-1} S^\top \exp\left(-Sp + \Phi\right) w\right\| \geq \|w\|,\end{aligned} \tag{3.76}$$

which concludes the proof that $D_\phi p$ is bounded by 1 and $F_{m,n}$ is a Lipschitz continuous map for any given m, n with Lipschitz constant β. Since $\beta < 1$, $F_{m,n}$ is a contraction mapping and it admits a fixed point.

Step 2: Let $\bar{U}, \bar{V}$ be a fixed point to $F_{n,m}$. Then since $(n, m) \to F_{n,m}$ is continuous, $(n, m) \to \left(\bar{U}(n, m), \bar{V}(n, m)\right)$ is also continuous in n, m. Besides the latter map is defined on a convex compact subset of $\mathbb{R}^{X+Y}$ since $\sum_x n_x + \sum_y m_y = M$.

Define map $G : (n, m) \to (n, m)$ as

$$\begin{aligned}G_x(n, m) &= \sum_{xy\in\mathcal{X}_0\,\mathcal{Y}} P_{x'|xy}\mu_{xy}(\bar{U}, \bar{V}, m, n) \\ G_y(n, m) &= \sum_{xy\in\mathcal{X}\,\mathcal{Y}_0} Q_{y'|xy}\mu_{xy}(\bar{U}, \bar{V}, n, m).\end{aligned} \tag{3.77}$$

By continuity of $(n, m) \to \left(\bar{U}(n, m), \bar{V}(n, m)\right)$, $(n, m) \to \mu(\bar{U}, \bar{V}, m, n)$ is continuous, hence G is continuous, and by Brouwer's theorem it admits a fixed point.

Conclusion

This dissertation aims at understanding returns to education and experience on labor markets, and understand how they are impacted by supply and demand changes. The first chapter evidences flattening returns to experience for high educated workers in France: average yearly entry wage growth on the medium-term is 11.9% for the 1998 cohort of new labor market entrants, and only 8.6% for the 2010 cohort. It also evidences the high heterogeneity of this flattening across occupations: some have benefitted from increased returns to an additional year of experience, while other have suffered a decline. More specifically, the occupations whose total share of the working population has increased are the ones that experience the most severe decline. This finding suggests an over-supply of workers leads to slower wage progression. I explore to mechanisms that could cause this effect: access to managerial positions and skill-occupation mismatch. I find promotions to managerial positions are increasingly uncommon among the 2010 cohort, compared to the 1998 cohort. I also find that although initial mismatch has stayed the same between 1998 and 2010, its weight in determining medium-term wage levels has grown between the two cohorts. The second chapter focuses on the high school wage premium in Portugal. It starts with the observation that the high school wage premium has decreased since the end of the 1980s, but to 50% for young workers. Two opposite interpretations could be given for this fact: trade or skill-biased technological change overtaken by an education expansion. I find that high school educated workers productivity has increased in all sectors of the Portuguese economy between the 1980s and today, and that the contraction in high school wage premium is driven by the rise in the relative numbers of high school educated workers. Finally the third chapter explores dynamic matching games, and shows that under reasonable assumptions, the equilibrium can be characterized and computed through various methods. This chapter provides a backbone for further empirical investigations in the dynamics of labor markets.

I plan on making this dissertation the basis for a broader research agenda on returns to education, experience, and wage inequality. My goal is to keep expanding on the tools from the matching literature to analyze labor markets through the lens of supply and demand equilibrium.

There are various issues I would like to explore next. First is the gender wage gap. A growing literature (Morchio and Moser (2019)) highlight the role of gender's differentiated valuation for amenities offered by firms to explain the wage discrepancy between men and women. An important question I would like to investigate is amenity valuation interacts with education, experience, and firm production, and how much it can account for the gender pay gap. Another interesting question is the role of labor market institutions in wage inequality: how are supply and demand equilibrium affected by the minimum wage level, severance policies or union wage bargaining? Third, I plan on exploring the underlying role of unobserved worker and firm productivity in the impact of education expansions on wage levels. On the worker side, productivity varies depending on latent ability (Abowd et al. (1999), Bonhomme et al. (2019)), which might in turn affect returns to education. If they are aware of their ability, workers will sort into education depending on its return. This is likely to change the distribution of underlying ability among uneducated and educated workers in times of education expansion, which in turns impacts wage levels. On the firm side, there may be both individual firm and co-worker effects on productivity. Accounting for unobserved heterogeneity is therefore key to understand why observed wage distributions evolve over time.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A Distributional Framework for Matched Employer Employee Data. *Econometrica*, 87(3):699–739.

Morchio, I. and Moser, C. (2019). The Gender Gap: Micro Sources and Macro Consequences. Technical Report 143, Society for Economic Dynamics. Publication Title: 2019 Meeting Papers.

Les Rendements de l'Éducation et de l'Expérience sur le Marché du Travail

Offre, Demande, et Appariement

Pauline Corblet

Résumé

Résumé

Les marchés du travail sont un élément incontournable des sociétés occidentales du 21^{e} siècle. Chaque individu en Europe et aux États-Unis a, à un moment donné de sa vie, une expérience sur ces marchés, soit en cherchant un emploi, soit en étant employé. Pas un jour ne passe sans que les médias et les gouvernements examinent et dissèquent le taux de chômage et les niveaux de salaire, considérés comme des mesures capitales de la santé de l'économie. Dans le domaine universitaire, une grande partie de la recherche économique est également consacrée aux marchés du travail. Les questions relatives à l'impact du commerce international et de la mondialisation, au niveau d'éducation de la main-d'œuvre, à la détermination des salaires et au rôle des institutions (entre autres) offrent des espaces de recherche dynamiques.

Cette thèse vise à comprendre les marchés du travail d'un point de vue microéconomique : elle modélise les choix des individus et des entreprises sur le marché de l'emploi par une maximisation de l'utilité (pour les individus) ou des profits (pour les entreprises), ce qui, au niveau agrégé, se traduit par une offre et une demande de main-d'œuvre. Elle utilise ensuite ce cadre pour décomposer les facteurs de détermination des salaires. Cette thèse s'inscrit donc dans la tradition néoclassique, dans laquelle les salaires sont des quantités d'équilibre déterminées par l'offre et la demande. Cependant, elle s'écarte de la théorie de base telle que présentée par Hicks (1932) : en tenant compte des idiosyncrasies individuelles, en introduisant des individus et des entreprises hétérogènes et en reconnaissant une concurrence imparfaite sur les marchés du travail, on tient compte ici d'un cadre plus riche que la vision standard qui postule que les salaires sont simplement égaux à la productivité marginale des employés.

Dans cette thèse, j'ai choisi de me concentrer sur un aspect crucial des marchés de l'emploi, à savoir les rendements salariaux de l'éducation et de l'expérience. On observe généralement que les individus diplômés du secondaire ou de l'université touchent un salaire plus élevé que les autres. Les individus ayant plus d'années d'expérience sur le marché du

travail sont aussi souvent mieux payés que leurs pairs moins expérimentés. Les rendements de l'éducation et de l'expérience se peuvent se combiner, en faveur des individus diplômés et expérimentés. Cependant, les écarts salariaux entre les travailleurs diplômés ou non (ou moins diplômés) et les travailleurs expérimentés et inexpérimentés varient dans le temps et dans l'espace. Comprendre pourquoi ces écarts existent en premier lieu et pourquoi ils diffèrent selon la décennie ou le pays est essentiel pour deux raisons. D'abord parce qu'ils sont des vecteurs importants d'inégalités : en 2016, les salaires représentaient 55 % du revenu total des ménages en France, et 72 % aux États-Unis (Rani and Furrer (2016)). Puisque les salaires sont la principale source de revenu des ménages dans la plupart des pays, l'inégalité des salaires entre les individus se traduit par une inégalité des revenus. [1] (Autor (2014)). La deuxième raison pour laquelle il est nécessaire de comprendre les écarts salariaux est qu'ils nous informent sur le fonctionnement des entreprises. En effet, la rémunération que les entreprises acceptent de verser à leurs employés dépend de leurs besoins : toutes choses égales par ailleurs, une entreprise est prête à rémunérer beaucoup plus un travailleur qui possède un ensemble de compétences essentielles à son fonctionnement qu'un travailleur qui ne les possède pas. L'augmentation de la prime à l'éducation (c'est-à-dire l'écart salarial moyen entre les travailleurs diplômés de l'université et les autres) entre la fin des années 1970 et le début des années 2000 aux États-Unis, au Royaume-Uni et au Canada (Krueger et al. (2010)), est interprétée de manière convaincante par une vaste littérature comme la conséquence du *skill-biased technological change*, selon le terme anglophone, c'est-à-dire un changement technologique en faveur des plus diplômés. Ainsi, la prime à l'éducation augmenterait en raison d'un changement dans la structure de production des entreprises qui accroît la productivité des individus diplômés par rapport aux non-diplômés (Katz and Murphy (1992), Acemoglu (1998), Goldin and Katz (2008), Autor et al. (2020)). Récemment, la tendance de la prime à l'éducation s'est aplatie, et elle a même diminué dans certains pays européens (par exemple en Allemagne, en Italie, en Espagne). Ce changement est révélateur de vastes transformations sur les marchés du travail où il s'est produit, et pourrait être dû à un effet de demande, à travers un retournement du changement technologique, soit à une modification de l'offre : tant l'Europe que les États-Unis ont connu d'importantes expansions éducatives au cours des 50 dernières années. Par conséquent, même si la demande des entreprises pour les employés diplômés a persisté, leur plus grand nombre pourrait être à l'origine de la baisse du salaire relatif. Comprendre comment l'offre et la demande interviennent dans la détermination des écarts salariaux est capital dans l'élaboration des politiques économiques et devrait façonner

[1] C'est particulièrement vrai dans le premier 90^{e} centile de la distribution des revenus : les revenus du capital sont concentrés au sommet de la distribution des revenus et constituent le principal moteur de l'inégalité entre les 90% les plus pauvres et les 10% les plus riches (Krueger et al. (2010)).

les politiques d'éducation et de production [2].

Distinguer les effets de l'offre et de la demande sur les niveaux de salaire est au cœur de cette thèse. Pour ce faire, elle développe un ensemble d'outils de modélisation des données qui empruntent à diverses branches de la littérature économique. Dans le premier chapitre, j'utilise des méthodes empiriques standards (la régression par moindres carrés ordinaires) pour évaluer l'impact moyen d'une ou plusieurs variables explicatives sur une variable de résultat. Dans le deuxième et troisième chapitre, je m'appuie sur la littérature en économétrie structurelle pour construire un modèle d'appariement entre les individus et les entreprises sur le marché du travail. Le modèle peut être soit statique, comme dans le deuxième chapitre, soit dynamique, comme dans le troisième. Dans les deux cas, il intègre les individus qui maximisent leur utilité et les entreprises qui maximisent leur profit, les salaires agissant comme des transferts d'équilibre. Il est important de noter que les agents du modèle sont hétérogènes dans de multiples dimensions : éducation, âge ou profession pour les individus, secteur ou emplacement pour les entreprises. Comme les individus perçoivent des niveaux d'utilité (hors salaire) différents en fonction de l'entreprise et que tous ne sont pas également productifs en fonction de leur employeur, le modèle génère de riches distributions de salaires et d'appariement. Enfin, j'exploite des données d'appariement entre employés et employeurs qui contiennent des informations précises sur les profils des individus et entreprises, ainsi que sur les salaires versés. À l'aide de ces données, je suis en mesure d'estimer structurellement les modèles susmentionnés en utilisant l'appariement et les salaires observés, afin d'estimer les paramètres gouvernant l'utilité des individus et le profit des entreprises sur le marché de l'emploi.

Le premier chapitre explore empiriquement les mécanismes de carrière qui se traduisent par des progressions salariales différentes par cohorte, le deuxième chapitre utilise des méthodes structurelles, à savoir les modèles d'appariement, particulièrement adaptés à l'étude des effets d'offre et de demande au niveau agrégé, pour évaluer l'effet simultané d'une expansion de l'éducation et d'un changement technologique sur les salaires. Le troisième chapitre développe un modèle d'appariement dynamique pour mesurer les attentes des agents sur les rendements futurs de leur appariement présent. Un bref résumé de chacun des chapitres de la thèse suit.

[2] Par exemple, le programme du lycée en France a été profondément réformé en 2019, entraînant une spécialisation renforcée des élèves. Il reste à évaluer si cela constitue un atout pour les nouveaux diplômés sur le marché du travail.

Chapitre 1. Le premier chapitre documente l'aplatissement des rendements de l'expérience pour les diplômés du supérieur en France entre 1998 et 2017. Je compare l'évolution des salaires à sept ans entre trois cohortes, ou générations, sorties du système scolaire ou supérieur en 1998, 2004 et 2010. Je documente la croissance moyenne des salaires par cohorte et par niveau d'éducation, et constate que pour les individus peu diplômés (n'ayant pas terminé le lycée ou diplômés du baccalauréat) le profil d'évolution du salaire moyen ne varie pas entre les générations. À l'inverse, les diplômés de l'enseignement supérieur profitent d'une croissance de salaire moyen plus forte en début de carrière, mais cette croissance s'aplatit entre la cohorte de 1998 et la cohorte de 2010.

L'économie et le système éducatif français ont connu des changements importants au cours de la période 1998-2017. Les cohortes 1998, 2004 et 2010 entrent donc sur le marché du travail dans des conditions sensiblement différentes : la cohorte 1998 est confrontée à un chômage élevé (supérieur à 10%), mais compte relativement peu de diplômés du supérieur. La cohorte 2004 bénéficie d'un faible taux de chômage et d'une forte demande des entreprises, mais compte davantage de diplômés de l'enseignement supérieur. Enfin, la cohorte de 2010 entre sur le marché du travail au milieu de la Grande Récession, et fait face à un fort taux de chômage et à une faible croissance. Parce que l'expansion de l'éducation française est encore forte dans les années 2000, encouragée par la création de licences professionnelles et la mise en œuvre du processus de Bologne, la cohorte 2010 compte nettement plus de diplômés du supérieur que ses prédécesseurs. Cela risque de nuire à leurs perspectives sur le marché du travail à plusieurs niveaux : premièrement, les jeunes de la génération 2010 sont confrontés à une baisse de la demande des entreprises. Ensuite, les diplômés du supérieur sont plus nombreux qu'auparavant. L'impact de ce second élément peut être pensé sous plusieurs angles : (Gaini et al. (2013) ; Dupray and Moullet (2010)). Premièrement, si le diplôme est un signal de productivité, l'augmentation du nombre de diplômés implique alors une diminution de leur productivité individuelle moyenne, qui peut se traduire par une progression plus lente des salaires. Une seconde approche considère le diplôme comme un moyen d'acquérir du capital humain. Les rendements de l'éducation repose alors sur cette acquisition. La diversification du système d'enseignement supérieur français, en modifiant le contenu des diplômes, a pu avoir un impact négatif sur l'acquisition du capital humain des jeunes diplômés. Enfin, une troisième approche basée sur le modèle néoclassique standard prédit une baisse du salaire des jeunes diplômés si leur nombre augmente simplement parce que le salaire est égal au produit marginal d'un travailleur supplémentaire. Si les entreprises produisent avec des rendements d'échelle décroissant, chaque travailleur supplémentaire fait diminuer le salaire moyen.

Ce chapitre s'attache donc à étudier empiriquement les raisons de la croissance différenciée

des salaires moyens en France depuis la fin des années 1990. Pour cela, j'utilise les enquêtes 'Générations' mises à disposition par le CEREQ (Centre d'Études et de Recherche sur les Qualifications). Les enquêtes sont présentées sous forme de données de panel et couvrent la vie active des sortants du système scolaire ou supérieur en 1998, 2004 et 2010 pendant sept ans, et fournissent une vision globale de l'insertion des jeunes sur le marché du travail français. Je décompose d'abord les différences de croissance des salaires moyens par catégorie socioprofessionnelle (PCS) en une marge extensive et une marge intensive. La marge extensive résulte d'un effet de composition dû à une évolution de la part représentée par chaque PCS entre les cohortes. La marge intensive repose sur la variation de la croissance annuelle des salaires par PCS entre cohortes. Les PCS qui affichent une marge intensive négative sont également celles pour lesquelles la marge extensive est la plus grande. En effet, les PCS qui accusent le ralentissement le plus important de la progression des salaires sont également celles qui connaissent le plus grand afflux de diplômés entre 1998 et 2010. Cette observation est conforme à une interprétation du ralentissement de la croissance des salaires en termes d'offre et de demande, selon laquelle une offre excédentaire de nouveaux diplômés les empêche d'atteindre les niveaux de salaire de leurs prédécesseurs. Une telle interprétation suggère d'explorer les mécanismes par lesquels une augmentation de l'offre de diplômés, conjuguée à une stagnation de la demande, influe sur la dynamique des salaires en début de carrière. Ce chapitre en explore deux : la promotion à des postes de d'encadrement d'une équipe, ou de manager, et l'inadéquation entre spécialité du diplôme et PCS. Les enquêtes 'Générations' montrent que l'obtention d'un poste de manager s'accompagne d'une augmentation de salaire à moyen terme. Ainsi, une diminution de la probabilité d'obtenir un tel poste ralentit la progression globale des salaires. Ceci est cohérent avec les conclusions de Kwon et al. (2010). J'examine ensuite l'argument de Liu et al. (2016), qui montrent qu'aux États-Unis, les diplômés de l'université ont souffert d'une adéquation diplôme-industrie dégradée pendant la Grande Récession, ce qui a conduit à des niveaux de salaire constamment inférieurs à ceux de leurs pairs plus âgés. En France, je n'observe pas d'aggravation de l'inadéquation (définie comme le niveau de salaire moyen en première année d'une spécialité de diplôme donnée au sein d'une PCS) entre 1998 et 2010 pour les diplômés du supérieur, mais je constate que son importance dans la détermination des salaires futurs s'est accrue entre Générations 1998 et 2010.

Chapitre 2. Entre les années 1970 et aujourd'hui, le nombre d'individu diplômés a fortement augmenté dans de nombreux pays. En conséquence, le rapport entre effectifs des individus diplômés et les effectifs des non-diplômées a augmenté sur le marché du travail. Ce chapitre cherche à comprendre comment cette augmentation a impacté l'appariement des individus

et des entreprises, c'est-à-dire quel type d'entreprise embauche quel type d'individu, ainsi que son impact sur les salaires, en utilisant un nouveau modèle d'appariement sur le marché du travail. Le modèle est structurellement estimé sur des données portugaises appariées employeurs-employés. Ce faisant, je suis en mesure de quantifier l'impact des changements de l'offre et de la demande sur l'allocation travailleur-entreprise et la structure des salaires.

Pour capturer les mécanismes d'offre (des individus sur le marché de l'emploi) et de demande (des entreprises) sur le marché du travail, ce chapitre construit un modèle d'appariement statique dans lequel une entreprise s'apparie avec plusieurs employés, et leur transfère un salaire (modèle avec utilité transférable). Les individus et les entreprises diffèrent de par leurs caractéristiques observées, qui sont résumées par un type multidimensionnel, ainsi que par un choc stochastique qui rend compte de l'hétérogénéité non observée. Une seule entreprise embauche plusieurs travailleurs, qui forment sa main-d'œuvre. Le surplus créé par l'appariement dépend des caractéristiques observables des entreprises ainsi que de celles de la main-d'œuvre. L'utilité est transférable sous la forme de salaires versés par l'entreprise à ses employés. Les entreprises cherchent à maximiser le profit total, qui est constitué de la production de la main-d'œuvre plus un choc aléatoire, moins les salaires versés. Les travailleurs maximisent leur utilité, qui se compose du salaire, plus une part non monétaire de préférence pour le type de l'entreprise et un choc aléatoire. À l'équilibre, les salaires équilibrent le marché et chaque agent s'apparie à sa meilleure option compte tenu des salaires. Le modèle peut générer une riche distribution des salaires qui dépend à la fois des caractéristiques observables des individus et de l'entreprise, ainsi que de la main-d'œuvre employée. Il prédit également un appariement d'équilibre, qui est la distribution conjointe des entreprises et de la main-d'œuvre. En utilisant à la fois l'appariement et les salaires, le modèle identifie séparément la production de l'entreprise et l'utilité non monétaire perçue par les individus.

Le cadre développé dans ce chapitre offre plus de flexibilité dans l'estimation que les modèles classiques d'offre et de demande de Katz and Murphy (1992) et Card and Lemieux (2001) : il identifie les préférences des individus en plus de la production de l'entreprise, et permet aux paramètres de varier dans le temps. En effet, en modélisant explicitement les choix d'appariement des entreprises et des individus, on peut utiliser à la fois l'appariement observé et les salaires observés, ce qui augmente les possibilités d'identification. Le modèle est estimé sur les données en supposant des fonctions paramétriques pour la production de l'entreprise et l'utilité non monétaire des individus. Les individus sont classés en deux niveaux d'éducation, les diplômés du secondaire et les non-diplômés, et en trois groupes d'âge, jeunes (16-34 ans), moyen (35-54 ans) et seniors (plus de 55 ans). Les entreprises se différencient par leur secteur d'activité. Conformément à la littérature, la production est une fonction imbriquée d'élasticité constante de substitution (CES), avec des paramètres de pro-

ductivité pour chaque niveau d'éducation qui varient d'un secteur à l'autre. Je suppose que les préférences des individus pour les entreprises dépendent de l'âge, du niveau d'éducation et du secteur de l'entreprise. En utilisant des prédictions du modèle pour l'appariement et les salaires, on l'estime structurellement sur des données appariées employeur-employé par maximum de vraisemblance sur la distribution conjointe de l'appariement et des salaires, séparément tous les trois ans.

Le modèle développé dans ce chapitre est lié à la fois aux problèmes d'affectation de plusieurs biens à un agent étudiés dans la littérature de 'market design', ou conception de marché, (Bikhchandani and Ostroy (2002), Vohra (2011)), et aux modèles d'appariement utilisés en économie de la famille (Choo and Siow (2006)). Le modèle développé dans ce chapitre comble l'écart entre ces deux littératures : il étend les affectations unilatérales à l'appariement bilatéral et généralise l'appariement entre deux agents à un appariement entre plus de deux agents. De plus, le cadre économétrique de Choo and Siow (2006) et Galichon and Salanié (2021) est étendu aux modèles d'appariement multiple.

Ce chapitre utilise le nouveau cadre théorique développé pour étudier le marché du travail portugais entre 1987 et 2017. Trois faits sont mis en évidence : premièrement, le pays opère une vaste expansion de sa population éduquée au cours de la période, ce qui se traduit par une augmentation spectaculaire de l'offre de diplômés du secondaire par rapport aux non-diplômés sur le marché du travail. Deuxièmement, le rendement de l'éducation diminue au cours de la période. Le rendement de l'éducation est défini comme l'écart salarial moyen entre les individus ayant obtenu leur diplôme d'études secondaires et ceux ne l'ayant pas obtenu. La diminution du rendement de l'éducation est particulièrement marquée chez les jeunes. Deux interprétations s'opposent pour expliquer ce fait. Premièrement, la diminution des rendements de l'éducation pourrait être la conséquence d'un effet commercial : le Portugal a rejoint l'Union européenne en 1986, et parce que le pays compte relativement plus de travailleurs non éduquées (qui n'ont pas fréquenté l'école secondaire) dans sa population active que les autres pays de l'UE, un modèle à la Heckscher-Ohlin prédit une augmentation des exportations de biens dont la production nécessite une main-d'œuvre non éduquée. La demande relative de main-d'œuvre non diplômée par rapport à la main-d'œuvre diplômée augmente et les rendements de l'éducation diminuent. La seconde interprétation repose sur un effet d'offre : même si la demande relative de main-d'œuvre diplômée par rapport à la main-d'œuvre non diplômée augmente, la formidable expansion de la population diplômée qui s'est produite au Portugal dans les années 1990 et 2000 pourrait diminuer la productivité marginale des travailleurs diplômés et entraîner une baisse des rendements de l'éducation. Le modèle décrit ci-dessus est en mesure de distinguer les deux interprétations possibles. Enfin, on observe que la répartition des diplômés du secondaire par rapport aux non-diplômés

entre les secteurs d'activité des entreprises devient de plus en plus déséquilibrée, en faveur des secteurs de services, transports et communications, qui emploient une part croissante de diplômés du secondaire. Les deux premiers faits impliquent que l'offre relative de diplômés du secondaire par rapport aux non-diplômés a augmenté plus rapidement que la demande relative des entreprises pour les diplômés du secondaire par rapport aux non-diplômés. Ce dernier suggère que l'appariement entre salariés et entreprises a évolué sur la période : soit parce que les entreprises des services, transports et communications nécessitent une part croissante de diplômés du secondaire, soit parce que la préférence des diplômés du secondaire pour ces entreprises se renforce.

Après estimation du modèle d'appariement, je constate que la demande relative de diplômés du secondaire dans les entreprises des secteurs des services, de l'industrie manufacturière et des transports et communications a considérablement augmenté au cours de la période, en particulier depuis le début des années 2010. Ce résultat est conforme à l'hypothèse de changement technologique biaisé en faveur des plus diplômées, plutôt qu'à un effet du commerce avec le reste de l'UE : il suggère que l'augmentation de la demande relative de main-d'œuvre diplômée par rapport à la main-d'œuvre non diplômée, favorisée par le changement technologique, est contrebalancée par l'augmentation de l'offre relative. Je constate également que la préférence des diplômés du secondaire de moins de 55 ans pour les secteurs de service, d'industrie manufacturière et de transports et communication a diminué, tandis que la part des moins de 55 ans dans la production des entreprises augmente par rapport aux salariés plus âgés. En plus du cadre classique de l'offre et de la demande, le modèle offre deux mécanismes supplémentaires d'évolution des rendements de l'expérience. Premièrement, une baisse de l'utilité non monétaire des individus a pour effet d'augmenter les salaires. Deuxièmement, la demande relative des entreprises pour les différentes classes d'âge varie dans le temps. J'effectue plusieurs exercices contrefactuels pour évaluer les actions distinctes des changements dans la démographie des travailleurs (à la fois dans l'éducation et la répartition par âge), la composition par secteur des entreprises, les paramètres de production et l'utilité non monétaire des individus, sur l'appariement et les rendements de l'expérience. Il apparaît que les changements démographiques sont le principal moteur des changements d'appariement. Les changements dans la composition de l'industrie, la demande des entreprises et l'utilité non monétaire des individus ont un effet modeste. Les rendements de l'éducation par tranche d'âge et par branche d'activité sont affectés négativement par les changements dans la démographie des travailleurs et la composition des secteurs d'activité (c'est-à-dire l'augmentation de la part des services) et positivement par les variations de la demande des entreprises. Ainsi les changements de productivité relative en faveur des diplômés du secondaire ont fait augmenter les rendements

de l'éducation, mais ne peuvent pas compenser la forte augmentation de l'offre relative de diplômés par rapport aux non-diplômés.

Chapitre 3. *Co-écrit avec Jeremy Fox et Alfred Galichon.* Ce chapitre adopte une perspective différente des deux premiers : au lieu d'étudier l'appariement des individus et des entreprises dans un monde statique, il explore comment des considérations dynamiques influencent l'appariement. En effet, sur de nombreux marchés d'appariement, y compris le marché du travail mais aussi le marché matrimonial, les agents tiennent compte du fait que leur type peut évoluer dans le temps, soit de manière déterministe (par exemple, l'âge des individus), soit en fonction de l'agent ou la personne avec laquelle ils se sont appariés (si un employé travaille dans une profession donnée, il accumulera du capital humain dans cette profession). Pour comprendre comment ces considérations influencent les choix des partenaires ou de l'employeur/employé, nous développons un modèle d'appariement dynamique. Les agents ont des types individuels, tels que l'éducation et l'expérience pour les salariés, et l'industrie et la profession pour les emplois. Lorsqu'ils décident avec qui s'apparier, les agents tiennent compte des rendements futurs attendus qui découlent d'un changement de type. À son tour, ce changement de type affectera les retours des appariements futurs. À chaque période, le marché s'équilibre via les salaires.

L'objectif de ce chapitre est de développer un modèle prêt à l'emploi utile de jeux d'appariement répétés qui généralise les jeux d'appariement statiques à un cadre dynamique. Il diffère également des deux chapitres précédents parce que le modèle développé n'est pas seulement applicable aux questions de travail, mais pourrait également s'appliquer en l'économie de la famille ou à l'organisation industrielle. Nous introduisons également des chocs stochastiques, ou erreurs économétriques, pour prendre en compte l'hétérogénéité inobservée dans les données. Le jeu d'appariement répété avec des erreurs économétriques s'apparente à une combinaison de deux articles de référence dans la littérature : Choo and Siow (2006) proposent un estimateur pour les jeux d'appariement statiques avec des erreurs logistiques, et Rust (1987) propose un estimateur pour les modèles de choix discrets dynamiques à agent unique, utilisant aussi des erreurs logistiques. Dans ce chapitre, nous combinons les deux pour obtenir un estimateur pour les jeux d'appariement dynamiques.

Dans notre cadre d'appariement répété, chaque agent a une variable d'état, ou type. S'apparier, ou rester sans partenaire, peut affecter l'évolution de cette variable d'état. À chaque période, les agents participent à un marché d'appariement avec des prix ou des transferts pour différents appariements. Compte tenu des prix d'équilibre du marché, chaque agent sélectionne le meilleur partenaire en prenant en compte son utilité future en fonction d'éventuel changement de sa variable d'état. La période suivante, le marché correspon-

dant rouvre, et de nouvelles correspondances se forment selon les nouveaux prix. Un jeu d'appariement répété peut avoir à la fois une dynamique individuelle et agrégée. Au niveau individuel, chaque agent résout un problème de programmation dynamique à agent unique, où chaque période, l'action de l'agent consiste à choisir un partenaire avec qui s'apparier. Au niveau agrégé, la variable d'état du marché correspondant est l'ensemble actuel de types d'agents actifs, ou variables d'état. Cette variable d'état agrégée évolue avec les décisions des agents individuels et se résume par une équation de Bellman au niveau de l'économie, résolue par un planificateur social. Nous développons d'abord le modèle sans erreur économétrique, puis en tenant compte des préférences individuelles sous forme de choc économétrique. Dans les deux cas, nous explorons deux méthodes différentes pour calculer l'équilibre agrégé : une méthode repose sur l'itération de la fonction de valeur du planificateur social sur une grille pour calculer sa valeur, et l'équilibre associé sur chaque point de la grille, et la seconde méthode utilise des réseaux de neurones pour minimiser une fonction de perte qui approche la recherche de point fixe de l'équation de Bellman du planificateur social.

L'un de nos résultats théoriques les plus importants est qu'il existe un équilibre stationnaire, à la fois avec et sans chocs économétriques : il existe une masse de variables d'état d'agent telle que, après que les appariements optimaux ont été choisis par les agents, les mêmes masses de type d'agents sont présentes à la période suivante. L'existence d'un équilibre stationnaire ne dépend pas des paramètres du modèle et permet éventuellement au chercheur d'ignorer la dynamique agrégée en imposant que le jeu d'appariement soit à un équilibre stationnaire. En nous concentrant sur l'équilibre stationnaire, nous introduisons encore deux autres méthodes pour le calculer : l'une résout le système d'équations de l'équilibre stationnaire en utilisant un solveur de programmation non linéaire. La deuxième méthode reformule le problème de la recherche d'un équilibre stationnaire comme un problème min-max et utilise l'algorithme primal-dual de Chambolle-Pock pour le résoudre. Nous montrons que ces deux méthodes peuvent s'adapter à des problèmes avec de nombreux types d'agents. En plus de calculer un équilibre stationnaire, nous pouvons étendre les mêmes estimateurs pour estimer structurellement les paramètres dans la production d'une correspondance avec un ensemble de données approprié.

Bibliographie

Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *The Quarterly Journal of Economics*, 113(4):1055–1089.

Autor, D., Goldin, C., and Katz, L. F. (2020). Extending the Race between Education and Technology. *AEA Papers and Proceedings*, 110:347–351.

Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the "other 99 percent". *Science*, 344(6186):843–851.

Bikhchandani, S. and Ostroy, J. M. (2002). The Package Assignment Model. *Journal of Economic Theory*, 107(2):377–406.

Choo, E. and Siow, A. (2006). Who Marries Whom and Why. *Journal of Political Economy*, 114(1):175–201.

Dupray, A. and Moullet, S. (2010). Offre éducative, valorisation des diplômes et effets de composition : deux générations de sortants de l'université au tournant des années 2000. *Revue d'economie politique*, Vol. 120(5):845–880.

Gaini, M., Leduc, A., and Vicard, A. (2013). Peut on parler de générations sacrifiées ? Entrer sur le marché du travail dans une période de mauvaise conjoncture économique. *Économie et Statistique*, n 462-463.

Galichon, A. and Salanié, B. (2021). Cupid's Invisible Hand: Social Surplus and Identification in Matching Models. *The Review of Economic Studies*.

Goldin, C. and Katz, L. (2008). *The Race Between Education and Technology.* Belknap Press for Harvard University Press.

Hicks, J. R. (1932). The Theory of Wages. *London Macmillan*.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Krueger, D., Perri, F., Pistaferri, L., and Violante, G. L. (2010). Cross-sectional facts for macroeconomists. *Review of Economic Dynamics*, 13(1):1–14.

Kwon, I., Milgrom, E. M., and Hwang, S. (2010). Cohort Effects in Promotions and Wages: Evidence from Sweden and the United States. *The Journal of Human Resources*, 45(3):772–808.

Liu, K., Salvanes, K. G., and Sørensen, E. O. (2016). Good skills in bad times: Cyclical skill mismatch and the long-term effects of graduating in a recession. *European Economic Review*, 84:3–17.

Rani, U. and Furrer, M. (2016). Decomposing income inequality into factor income components: Evidence from selected G20 countries. ILO Research Paper No. 15, International Labor Office.

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, 55(5):999–1033.

Vohra, R. V. (2011). *Mechanism Design: A Linear Programming Approach.* Econometric Society Monographs. Cambridge University Press, Cambridge.